Planck 2015 results

Astronomy & Astrophysics

Special feature

Planck 2015 results

XII. Full focal plane simulations

Planck Collaboration: P. A. R. Ade[87], N. Aghanim[60], M. Arnaud[75], M. Ashdown[71,5], J. Aumont[60], C. Baccigalupi[86], A. J. Banday[95,9], R. B. Barreiro[66], J. G. Bartlett[1,68], N. Bartolo[31,67], E. Battaner[97,98], K. Benabed[61,94], A. Benoît[58], A. Benoit-Lévy[25,61,94], J.-P. Bernard[95,9], M. Bersanelli[34,49], P. Bielewicz[84,9,86], J. J. Bock[68,11], A. Bonaldi[69], L. Bonavera[66], J. R. Bond[8], J. Borrill[14,90,*], F. R. Bouchet[61,89], F. Boulanger[60], M. Bucher[1], C. Burigana[48,32,50], R. C. Butler[48], E. Calabrese[92], J.-F. Cardoso[76,1,61], G. Castex[1], A. Catalano[77,74], A. Challinor[63,71,12], A. Chamballu[75,16,60], H. C. Chiang[28,6], P. R. Christensen[85,37], D. L. Clements[56], S. Colombi[61,94], L. P. L. Colombo[24,68], C. Combet[77], F. Couchot[73], A. Coulais[74], B. P. Crill[68,11], A. Curto[66,5,71], F. Cuttaia[48], L. Danese[86], R. D. Davies[69], R. J. Davis[69], P. de Bernardis[33], A. de Rosa[48], G. de Zotti[45,86], J. Delabrouille[1], J.-M. Delouis[61,94], F.-X. Désert[54], C. Dickinson[69], J. M. Diego[66], K. Dolag[96,81], H. Dole[60,59], S. Donzelli[49], O. Doré[68,11], M. Douspis[60], A. Ducout[61,56], X. Dupac[39], G. Efstathiou[63], F. Elsner[25,61,94], T. A. Enßlin[81], H. K. Eriksen[64], J. Fergusson[12], F. Finelli[48,50], O. Forni[95,9], M. Frailis[47], A. A. Fraisse[28], E. Franceschi[48], A. Frejsel[85], S. Galeotta[47], S. Galli[70], K. Ganga[1], T. Ghosh[60], M. Giard[95,9], Y. Giraud-Héraud[1], E. Gjerløw[64], J. González-Nuevo[20,66], K. M. Górski[68,99], S. Gratton[71,63], A. Gregorio[35,47,53], A. Gruppuso[48], J. E. Gudmundsson[28], F. K. Hansen[64], D. Hanson[82,68,8], D. L. Harrison[63,71], S. Henrot-Versillé[73], C. Hernández-Monteagudo[13,81], D. Herranz[66], S. R. Hildebrandt[68,11], E. Hivon[61,94], M. Hobson[5], W. A. Holmes[68], A. Hornstrup[17], W. Hovest[81], K. M. Huffenberger[26], G. Hurier[60], A. H. Jaffe[56], T. R. Jaffe[95,9], W. C. Jones[28], M. Juvela[27], A. Karakci[1], E. Keihänen[27], R. Keskitalo[14], K. Kiiveri[27,44], T. S. Kisner[79], R. Kneissl[38,7], J. Knoche[81], M. Kunz[18,60,2], H. Kurki-Suonio[27,44], G. Lagache[4,60], J.-M. Lamarre[74], A. Lasenby[5,71], M. Lattanzi[32], C. R. Lawrence[68], R. Leonardi[39], J. Lesgourgues[62,93], F. Levrier[74], M. Liguori[31,67], P. B. Lilje[64], M. Linden-Vørnle[17], V. Lindholm[27,44], M. López-Caniego[39,66], P. M. Lubin[29], J. F. Macías-Pérez[77], G. Maggio[47], D. Maino[34,49], N. Mandolesi[48,32], A. Mangilli[60,73], M. Maris[47], P. G. Martin[8], E. Martínez-González[66], S. Masi[33], S. Matarrese[31,67,42], P. McGehee[57], P. R. Meinhold[29], A. Melchiorri[33,51], J.-B. Melin[16], L. Mendes[39], A. Mennella[34,49], M. Migliaccio[63,71], S. Mitra[55,68], M.-A. Miville-Deschênes[60,8], A. Moneti[61], L. Montier[95,9], G. Morgante[48], D. Mortlock[56], A. Moss[88], D. Munshi[87], J. A. Murphy[83], P. Naselsky[85,37], F. Nati[28], P. Natoli[32,3,48], C. B. Netterfield[21], H. U. Nørgaard-Nielsen[17], F. Noviello[69], D. Novikov[80], I. Novikov[85,80], C. A. Oxborrow[17], F. Paci[86], L. Pagano[33,51], F. Pajot[60], D. Paoletti[48,50], F. Pasian[47], G. Patanchon[1], T. J. Pearson[11,57], O. Perdereau[73], L. Perotto[77], F. Perrotta[86], V. Pettorino[43], F. Piacentini[33], M. Piat[1], E. Pierpaoli[24], D. Pietrobon[68], S. Plaszczynski[73], E. Pointecouteau[95,9], G. Polenta[3,46], G. W. Pratt[75], G. Prézeau[11,68], S. Prunet[61,94], J.-L. Puget[60], J. P. Rachen[22,81], R. Rebolo[65,15,19], M. Reinecke[81], M. Remazeilles[69,60,1], C. Renault[77], A. Renzi[36,52], I. Ristorcelli[95,9], G. Rocha[68,11], M. Roman[1], C. Rosset[1], M. Rossetti[34,49], G. Roudier[1,74,68], J. A. Rubiño-Martín[65,19], B. Rusholme[57], M. Sandri[48], D. Santos[77], M. Savelainen[27,44], D. Scott[23], M. D. Seiffert[68,11], E. P. S. Shellard[12], L. D. Spencer[87], V. Stolyarov[5,91,72], R. Stompor[1], R. Sudiwala[87], D. Sutton[63,71], A.-S. Suur-Uski[27,44], J.-F. Sygnet[61], J. A. Tauber[40], L. Terenzi[41,48], L. Toffolatti[20,66,48], M. Tomasi[34,49], M. Tristram[73], M. Tucci[18], J. Tuovinen[10], L. Valenziano[48], J. Valiviita[27,44], B. Van Tent[78], P. Vielva[66], F. Villa[48], L. A. Wade[68], B. D. Wandelt[61,94,30], I. K. Wehus[68], N. Welikala[92], D. Yvon[16], A. Zacchei[47], and A. Zonca[29]

(Affiliations can be found after the references)

ABSTRACT

We present the 8th full focal plane simulation set (FFP8), deployed in support of the *Planck* 2015 results. FFP8 consists of 10 fiducial mission realizations reduced to 18 144 maps, together with the most massive suite of Monte Carlo realizations of instrument noise and CMB ever generated, comprising 10^4 mission realizations reduced to about 10^6 maps. The resulting maps incorporate the dominant instrumental, scanning, and data analysis effects, and the remaining subdominant effects will be included in future updates. Generated at a cost of some 25 million CPU-hours spread across multiple high-performance-computing (HPC) platforms, FFP8 is used to validate and verify analysis algorithms and their implementations, and to remove biases from and quantify uncertainties in the results of analyses of the real data.

1. Introduction

Planck[1] is the third satellite to study the cosmic microwave background (CMB) radiation. Launched in May 2009, *Planck* started its science observations from the L_2 Lagrange point in August 2009 and completed an all-sky survey (Planck Collaboration I 2014) approximately every six months until it was decommissioned in October 2013. *Planck* carried two instruments.

The High Frequency Instrument (HFI; Lamarre et al. 2010; Planck HFI Core Team 2011), comprising 52 detectors[2] at six frequencies (100, 143, 217, 353, 545, and 857 GHz), completed its observations when its cryogens were exhausted in January 2012 after almost five surveys, while the Low Frequency Instrument (LFI; Bersanelli et al. 2010; Mennella et al. 2011), comprising 22 detectors at three frequencies (30, 44, and 70 GHz), continued operating throughout the satellite lifetime, completing more than eight surveys.

The second release of *Planck* data (hereafter PR2-2015) is based on five HFI and eight LFI surveys (LFI Survey 9 having been reserved for atypical scans designed to assist in the

---

* Corresponding author: J. Borrill, e-mail: jdborrill@lbl.gov
[1] *Planck* (http://www.esa.int/Planck) is a project of the European Space Agency (ESA) with instruments provided by two scientific consortia funded by ESA member states and led by Principal Investigators from France and Italy, telescope reflectors provided through a collaboration between ESA and a scientific consortium led and funded by Denmark, and additional contributions from NASA (USA).

[2] The data from two detectors proved to be unusable, so only 50 detectors are included in the analysis.

Article published by EDP Sciences

understanding of systematic effects) and, for the first time, includes results in polarization as well as intensity. PR2-2015 is accompanied by a suite of papers, of which this is one, together with an online explanatory supplement, an ESA legacy data archive[3] and NASA partial mirror[4], as well as resources to access and manipulate the full simulation suite described here[5].

Simulations play a number of important roles in the analysis of these *Planck* data, including:

1. validating and verifying tools used to measure instrument characteristics by simulating data with known instrument characteristics, applying the tools used on the real data to measure these, and verifying the accuracy of their recovery;

2. quantifying systematic effect residuals by simulating data with some particular systematic effect included, applying the treatment used on the real data to mitigate that effect, and measuring the residuals;

3. validating and verifying data analysis algorithms and their implementations by simulating data with known science inputs (cosmology and foreground sky) and detector characteristics (beam, bandpass, and noise spectrum), applying the analyses used on the real data to extract this science, and verifying the accuracy of its recovery;

4. debiasing and quantifying uncertainties in the analysis of the real data by generating massive sets of Monte Carlo (MC) realizations of both the noise and the CMB and passing them separately through the analyses used on the real data to quantify biases and uncertainties.

The first two items are specific to each instrument, and distinct pipelines have been developed and employed by the LFI and HFI data processing centres (DPCs); details are provided in the DPC processing papers (Planck Collaboration II 2016; Planck Collaboration III 2016; Planck Collaboration IV 2016; Planck Collaboration V 2016; Planck Collaboration VI 2016; Planck Collaboration VII 2016; Planck Collaboration VIII 2016). However, the last two items require consistent simulations of LFI and HFI in tandem; such simulations are beyond the scope of either single-instrument pipeline. Furthermore, generating the Monte Carlo simulations is the most computationally intensive part of the *Planck* data analysis and requires computational capacity and capability far beyond those available at the DPCs. As a result, a massively parallel cross-instrument suite of codes has been developed to provide the collaboration with self-consistent simulations of all detectors at all frequencies – refered to as the full focal plane (FFP) – primarily using the high performance computing (HPC) resources at the National Energy Research Scientific Computing Center[6] (NERSC) in the USA and at CSC–IT Center for Science[7] (CSC) in Finland. Since its first deployment in 2006, this suite of codes has been used to generate a sequence of FFP simulation-sets of increasing veracity, complexity, and volume. The second *Planck* data release is supported by FFP8.

This paper is laid out as follows. Section 2 outlines the overall specification of FFP8, including an explicit listing of all of the maps to be generated, while Sect. 3 details the inputs to FFP8, comprising the instrument and observation models and the syntheic foreground and CMB skies. Section 4 describes the three major pipelines used to generate FFP8, both in terms of the algorithms and their massively parallel implementations, and Sect. 5

covers the execution of each of these pipelines at scale on specific HPC systems; the efficient use of scarce HPC resources is so critical to the generation of the FFP simulations that these sections address the key elements in achieving this in considerable detail. Finally Sect. 6 summarizes the results and current status.

2. Specifications

The FFP8 simulations include a set of fiducial mission realizations together with separate sets of Monte Carlo realizations of the CMB and the instrument noise. They contain the dominant instrumental (detector beam, bandpass, and correlated noise properties), scanning (pointing and flags), and analysis (mapmaking algorithm and implementation) effects.

In addition to the baseline maps made from the data from all detectors at a given frequency for the entire mission, there are a number of data cuts that are mapped both for systematics tests and to support cross-spectral analyses. These include:

– detector subsets ("detsets"), comprising the individual unpolarized detectors and the polarized detector quadruplets corresponding to each leading/trailing horn pair[8];

– mission subsets, comprising the surveys, years, and half-missions, with exact boundary definitions given in Planck Collaboration II (2016) and Planck Collaboration VII (2016) for LFI and HFI, respectively; and

– half-ring subsets, comprising the data from either the first or the second half of each pointing-period ring,

The various combinations of these data cuts then define 1134 maps, as enumerated in the top section of Table 1. The different types of map are then named according to their included detectors (channel or detset), interval (mission, half-mission, year or survey), and ring-content (full or half-ring); for example the baseline maps are described as channel/mission/full, etc.

The fiducial realizations include instrument noise (Sect. 3.1.2), astrophysical foregrounds (Sect. 3.2), and the lensed scalar, tensor, and non-Gaussian CMB components (Sect. 3.3), and are primarily designed to support the validation and verification of analysis codes. To test our ability to detect tensor modes and non-Gaussianity, we generate five CMB realizations with various cosmologically interesting – but undeclared – values of the tensor-to-scalar ratio r and non-Gaussianity parameter f_{NL} (Table 2). To investigate the impact of differences in the bandpasses of the detectors at any given frequency, the foreground sky is simulated using both the individual detector bandpasses and a common average bandpass, to include and exclude the effects of bandpass mismatch. To check that the PR2-2015 results are not sensitive to the exact cosmological parameters used in FFP8 we subsequently generated FFP8.1, exactly matching the PR2-2015 cosmology.

Since mapmaking is a linear operation, the easiest way to generate all of these different realizations is to build the full set of 1134 maps of each of six components:

– the lensed scalar CMB (`cmb_scl`);
– the tensor CMB (`cmb_ten`);
– the non-Gaussian complement CMB (`cmb_ngc`);
– the foregrounds including bandpass mismatch (`fg_bpm`);
– the foregrounds excluding bandpass mismatch (`fg_nobpm`);
– the noise (noise).

---

[8] The HFI DPC sometimes refers to entire channels as detset0; here detset only refers to subsets of detectors.

Table 1. Numbers of fiducial, MC noise, and MC CMB maps at each frequency by detector subset, data interval, and data cut.

			Frequency									Total
Detectors	Interval	Cut	30	44	70	100	143	217	353	545	857	
					FFP8 and FFP8.1 fiducial maps							
Channel	Mission	Full	1	1	1	1	1	1	1	1	1	9
		Half-ring	2	2	2	2	2	2	2	2	2	18
	Half-mission	Full	2	2	2	2	2	2	2	2	2	18
		Half-ring	4	4	4	4	4	4	4	4	4	36
	Year	Full	4	4	4	2	2	2	2	2	2	24
		Half-ring	8	8	8	4	4	4	4	4	4	48
	Survey	Full	8	8	8	4	4	4	4	4	4	48
		Half-ring	16	16	16	8	8	8	8	8	8	96
DetSet	Mission	Full	. . .	. . .	3	2	5	6	6	3	4	29
		Half-ring	. . .	. . .	6	4	10	12	12	6	8	58
	Half-mission	Full	. . .	. . .	6	4	10	12	12	6	8	58
		Half-ring	. . .	. . .	12	8	20	24	24	12	16	116
	Year	Full	. . .	. . .	12	4	10	12	12	6	8	64
		Half-ring	. . .	. . .	24	8	20	24	24	12	16	128
	Survey	Full	. . .	. . .	24	8	20	24	24	12	16	128
		Half-ring	. . .	. . .	48	16	40	48	48	24	32	256
Total			45	45	180	81	162	189	189	108	135	1134
					FFP8 MC noise maps							
Channel	Mission	Full	10 000	10 000	10 000	10 000	10 000	10 000	10 000	10 000	10 000	90 000
		Half-ring	20 000	20 000	20 000	200	200	200	200	200	200	61 200
	Half-mission	Full	2000	2000	2000	2000	2000	2000	2000	2000	2000	18 000
		Half-ring	4000	4000	4000	400	400	400	400	400	400	14 400
	Year	Full	4000	4000	4000	2000	2000	2000	2000	2000	2000	24 000
		Half-ring	8000	8000	8000	400	400	400	400	400	400	26 400
	Survey	Full	8000	8000	8000	4000	4000	4000	4000	4000	4000	48 000
		Half-ring	16 000	16 000	16 000	800	800	800	800	800	800	52 800
DetSet	Mission	Full	. . .	. . .	3000	2000	5000	6000	6000	3000	13 000	38 000
		Half-ring	. . .	. . .	6000	400	1000	1200	1200	600	800	11 200
	Half-mission	Full	. . .	. . .	1800	4000	10 000	12 000	12 000	6000	8000	53 800
		Half-ring	. . .	. . .	3600	800	2000	2400	2400	1200	1600	14 000
	Year	Full	. . .	. . .	3600	4000	10 000	12 000	12 000	6000	8000	55 600
		Half-ring	. . .	. . .	7200	800	2000	2400	2400	1200	1600	17 600
	Survey	Full	. . .	. . .	7200	8000	20 000	24 000	24 000	12 000	16 000	111 200
		Half-ring	. . .	. . .	14 400	1600	4000	4800	4800	2400	3200	35 200
Total			72 000	72 000	118 800	41 400	73 800	84 600	84 600	52 200	72 000	671 400
					FFP8 MC CMB maps							
Channel	Mission	Full	10 000	10 000	10 000	10 000	10 000	10 000	10 000	10 000	10 000	90 000
	Half-mission	Full	20 000	20 000	20 000	20 000	20 000	20 000	20 000	20 000	20 000	180 000
Detset	Mission	Full	. . .	. . .	30 000	20 000	50 000	40 000	20 000	20 000	10 000	190 000
Total			30 000	30 000	60 000	50 000	80 000	70 000	50 000	50 000	40 000	460 000
					FFP8.1 MC CMB maps							
Channel	Mission	Full	1000	1000	1000	1000	1000	1000	1000	1000	1000	9000
	Half-mission	Full	2000	2000	2000	2000	2000	2000	2000	2000	2000	18 000
Total			3000	3000	3000	3000	3000	3000	3000	3000	3000	27 000

Table 2. Values of tensor (r) and non-Gaussianity (f_{NL}) parameters for the FFP8 baseline and FFP8a–d blind challenge cases.

Data set	r	f_{NL}
FFP8	0	0
FFP8a	0.222	−4.311
FFP8b	0.046	8.590
FFP8c	0.088	7.140
FFP8d	0.153	−2.181

We then sum these, weighting the tensor and non-Gaussian complement maps with $\sqrt{r}$ and f_{NL}, respectively, and including one of the two foreground maps, to produce 10 total maps of each type. The complete fiducial data set then comprises 18 144 maps.

While the full set of maps can be generated for the fiducial cases, for the 10^4-realization MC sets this would result in some 10^7 maps and require about 6 PB of storage. Instead, therefore, the number of realizations generated for each type of map is chosen to balance the improved statistics it supports against the computational cost of its generation and storage. In practice, we generate 10^6 MC maps in total, including the full 10^4 realizations for all channel/mission/full maps for both noise and CMB. As detailed in Table 1, the remaining noise MCs sample broadly across all data cuts, while the additional CMB MCs are focused on the channel/half-mission/full maps and the subset of the detset/mission/full maps required by the `Commander` component separation code (Planck Collaboration X 2016).

3. Inputs

The inputs to an FFP simulation consist of the characterization of the mission (the satellite pointing and each detector's data flags), the intrument (the focal plane layout and the noise statistics, beam, and bandpass of each of the detectors), and the synthetic sky to be observed.

3.1. Mission and instrument characteristics

The goal of FFP8 is to simulate the *Planck* mission as accurately as possible; however, there are a number of known systematic effects that are not included, either because they are removed in the pre-processing of the time-ordered data (TOD), or because they are insufficiently well-characterized to simulate reliably, or because their inclusion (simulation and removal) would be too computationally expensive. These systematic effects are discussed in detail in Planck Collaboration II (2016) and Planck Collaboration VII (2016) and include:

- cosmic ray glitches (HFI);
- spurious spectral lines from the 4-K cooler electronics (HFI);
- nonlinearity in the analogue-to-digital converter (HFI);
- imperfect reconstruction of the focal plane geometry.

Note that if the residuals from the treatment of any of these effects could be mapped in isolation, then maps of such systematics could simply be added to the existing FFP8 maps to improve their correspondence to the real data.

3.1.1. Pointing and flags

The FFP8 detector pointing is calculated by interpolating the satellite attitude to the detector sample times and by applying a fixed rotation from the satellite frame into the detector frame. The fixed rotations are determined by the measured focal plane geometry (Planck Collaboration IV 2016; Planck Collaboration VII 2016), while the satellite attitude is
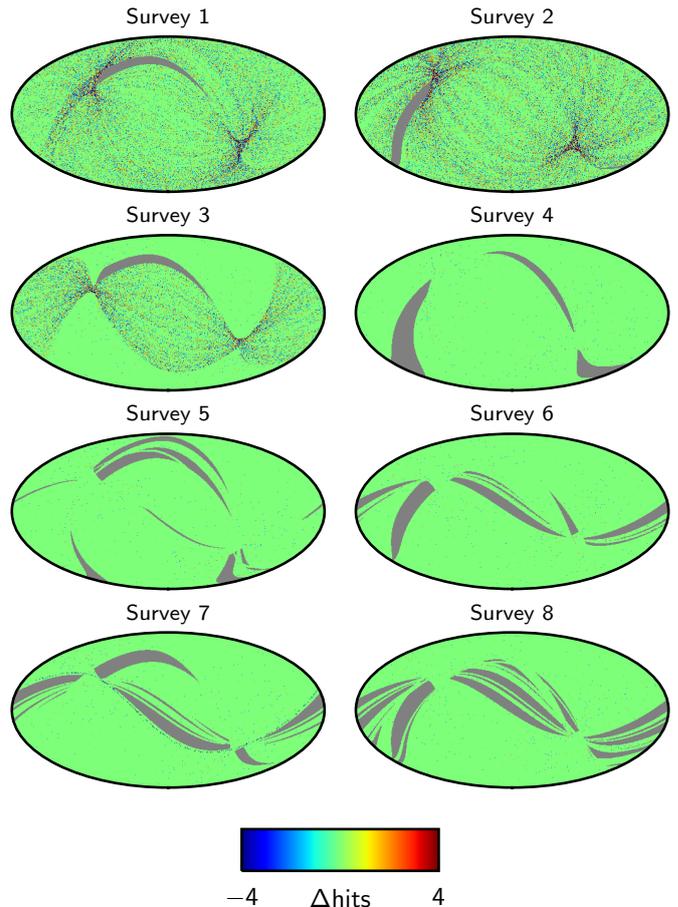
Fig. 1. Survey hit map differences between FFP8 (using `TOAST` pointing) and LFI at 70 GHz. The diffuse error pattern during Surveys 1–3 was caused by differences in pointing interpolation. The matching HFI comparison is presented in Fig. 2. The FFP pipeline rejected 2.5 min of integration time on Survey 7, leaving a small deficit of hits on a single scanning ring that is not visible in this projection. A hit is assigned to a pixel where the detector beam centre falls midway through a sample integration period.

described in the *Planck* attitude history files (AHF). The FFP pointing expansion reproduces the DPC pointing to sub-arcsecond accuracy, except for three short and isolated instances during Surveys 6–8 where the LFI sampling frequency was out of specification. Pixelization of the information causes the pointing error to be quantized to either zero (majority of cases) or the distance between pixel centres (3$'$.4 and 1$'$.7 for LFI and HFI, respectively). Since we need a single reconstruction that will serve both instruments efficiently in a massively parallel environment, we use the pointing provided by the Time Ordered Astrophysics Scalable Tools (`TOAST`) package.

In Fig. 1 we compare LFI DPC and `TOAST` "hit maps," which show the number of times a detector beam centre has fallen in a given pixel at the mid-point of a data sample, for the specified time period. There is a low level, diffuse, difference during the first three surveys, caused by slight differences in the way the pointing correction is interpolated in the two pipelines. The net effect is that some samples near pixel boundaries are assigned to different pixels in the different pipelines. Due to pixels being vastly oversampled, this affects about 25% of the sky pixels in the first three surveys, while the fraction of discrepant pixels after Survey 3 is about 0.5%. In the affected pixels, the relative difference in hit count is 0.2 ± 0.1% throughout. For Survey 1,

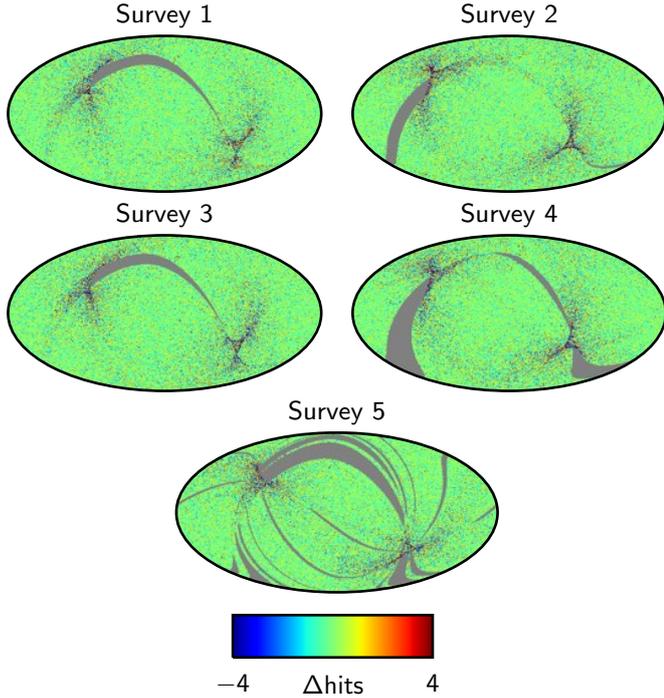**Fig. 3.** Example comparison of the input and recovered noise PSDs for one pointing period of a signal+noise simulation of HFI bolometer 100-1a. The shaded regions around the input model reflect the asymmetric realization scatter of the estimated PSDs at 68% and 95% confidence intervals in each of the 707 logarithmically-placed frequency bins.

the 70 GHz median hit count is 812 and 68% of the hit counts fall between 604 and 1278.

Except for Survey 7, the channel integration time differs by about a second, with the flight data maps always containing slightly more data. The larger difference during LFI Survey 7 occurs where TOAST rejected 2.5 min of data due to their irregular sampling rate. All of these differences are entirely negligible compared to the 14 Ms of overall integration time in each survey. For Survey 7 the relative difference in hits in differing pixels is slightly elevated to 0.3%, with an asymmetric standard deviation of 0.6%.

The survey-by-survey HFI comparison in Fig. 2 does not have a diffuse pattern matching Fig. 1 that ends during Survey 3, indicating a better agreement in the pointing correction interpolation between TOAST and HFI DPC. The more pronounced ecliptic pole caustics are caused by a truncation error in the TOD time stamps that were used as a basis of pointing expansion. In addition, the HFI's smaller pixel size and higher sampling frequency increases the likelihood that small numerical differences can result in samples being assigned across a pixel boundary. On average, about 52% of the pixels are affected and exhibit hit differences at the level of $0.5 \pm 0.3\%$. For 100 GHz and Survey 1, the median hit count is 224 and 68% of the hit counts fall between 164 and 361. The channel integration time between the two pipelines in the HFI case shows negligible differences, of the order of a few tenths of a second for all surveys.

3.1.2. Noise

We require simulated noise realizations that are representative of the noise in the flight data, including variations in the noise power spectral density (PSD) of each detector over time. To obtain these we developed a noise estimation pipeline
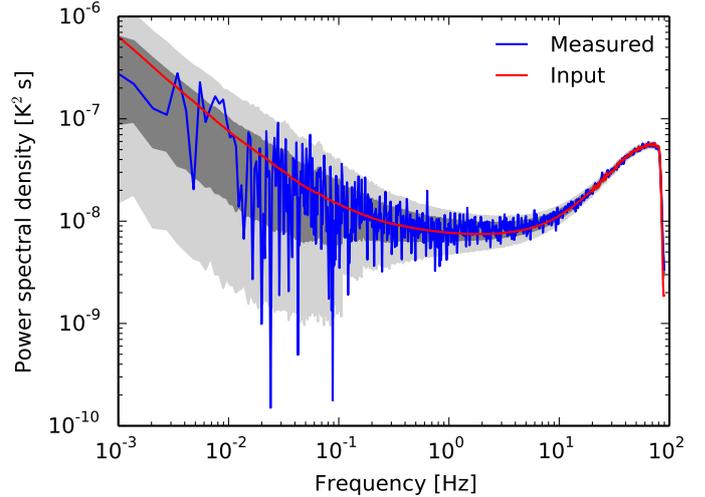
complementary to those of the DPCs. The goal of DPC noise estimation is to monitor instrument health and to derive optimal noise weighting, whereas our estimation is optimized to feed into noise simulation. Key features are the use of full mission maps for signal subtraction, long (about 24 h) realization length, and the use of autocorrelation functions in place of Fourier transforms to handle flagged and masked data (HFI).

The PSD for each detector and each roughly 45 min pointing period was estimated as follows.

1. Read in a 15% mask for the Galaxy and point sources.
2. Read the calibrated and dipole-subtracted time-ordered data and quality flags for the interval.
3. Expand detector pointing for the interval.
4. Scan and subtract a signal estimated from the full mission, full frequency map using the detector pointing weights.
5. Subtract the mean of the interval.
6. Then either
   – Fourier transform the signal-subtracted TOD for a noise PSD estimate, correcting for missing samples by dividing by the unflagged fraction (LFI),
   or
   – estimate the autocovariance function from the unflagged and unmasked samples and Fourier transform the autocovariance function into a noise PSD (HFI).
7. Bin the noise PSD estimate into 300 logarithmically spaced bins.
8. Average the noise estimates into approximately daily noise spectra.
9. Fit an analytic $1/f$ model to the measured spectra.

We show the results of running our pipeline on a signal+noise simulation in Fig. 3.

As a part of this process we observed significant low frequency (~1 mHz) correlated noise in the flight data between polarized HFI bolometers that share a horn[9], designated *a* and *b*.

---

[9] Each of the polarized horns on the *Planck* focal plane has two orthogonally polarized detectors. Due to the scanning strategy, it takes two polarized horns, i.e., a quadruplet of detectors, to solve for polarization away from the regions close to the ecliptic poles where the scanning rings intersect.
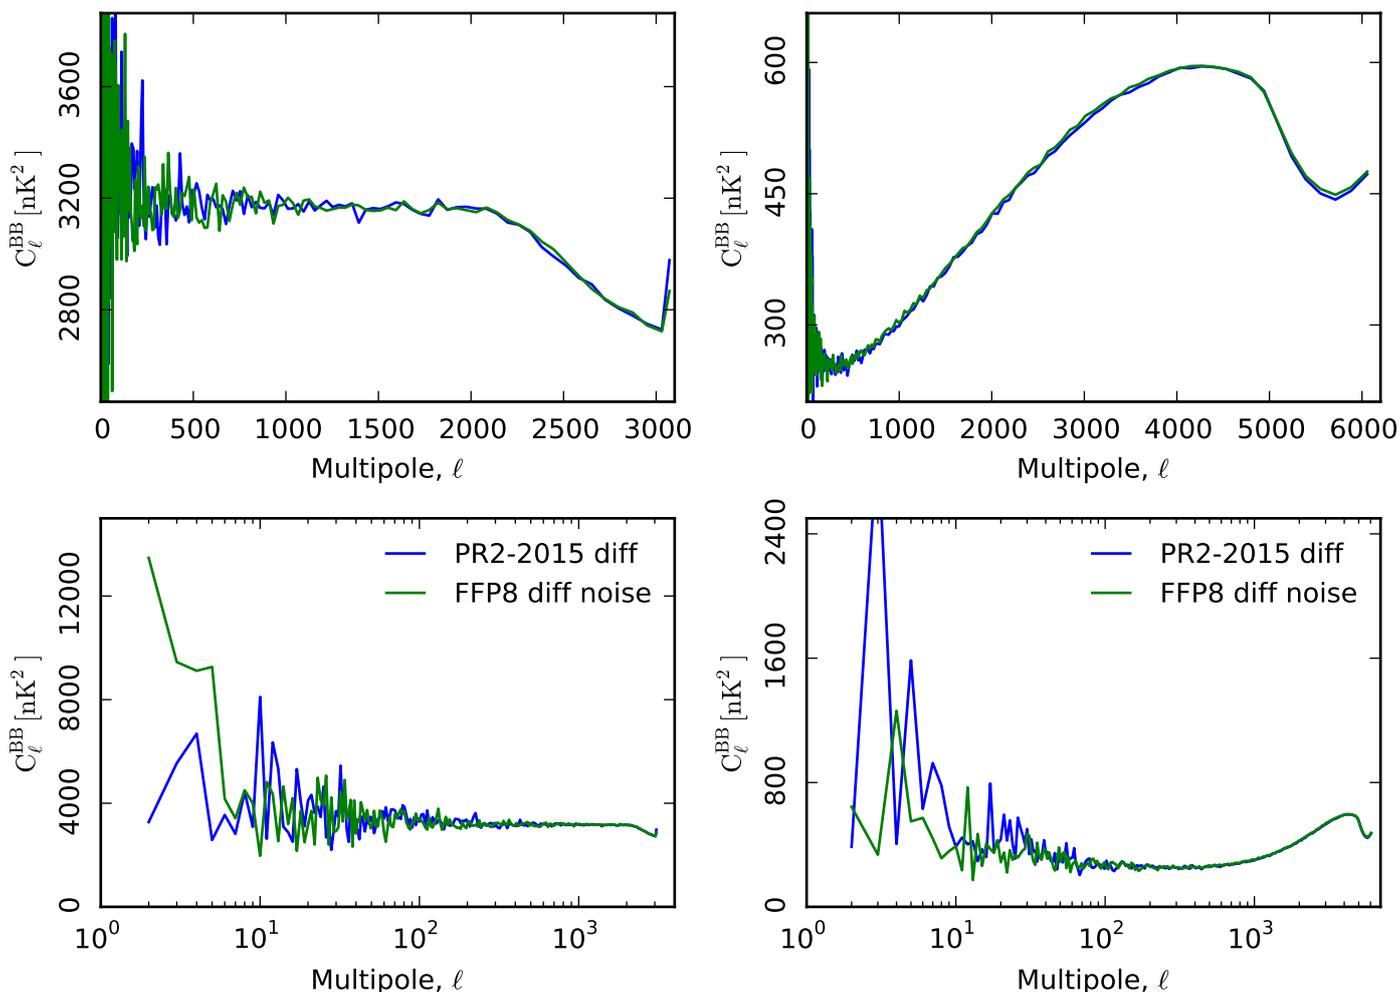
Fig. 4. *BB* spectra of 70 GHz (*left*) and 100 GHz (*right*) half-ring half difference noise maps, for both FFP8 simulations (green) and flight data (black). Differences at the low-ℓ end are caused by sample variance. These are pseudo-spectra, computed on 75% of the sky with the Galactic plane and point sources masked. *Top*: linear horizontal axis to show small-scale behaviour. *Bottom*: logarithmic horizontal axis to show large-scale behaviour.

These pairs of detectors form a polarization-sensitive unit, as their relative difference is sensitive to linear polarization of the incident radiation, and the existence of a correlated noise component between them is a known feature of the HFI intrument Planck Collaboration VI (2014). When this correlated noise was not included in the simulations, it left a clear signature in the mismatch between our noise simulations and the half ring difference maps. Specifically, a/b correlated noise adds coherently in the $(a + b)$ combination that is sensitive to temperature modulation of the signal, whereas it cancels in the $(a − b)$ combination that detects polarization modulation of the signal. Without simulating the correlated noise component, our simulated noise maps had too little large-scale noise in the *TT* spectrum, but excess noise in the *EE* and *BB* spectra. We estimated the correlated noise spectrum by differencing the $(a + b)$ and $(a − b)$ noise estimates, and added the correlated component into our noise simulation. FFP8 is the first simulation to include realistic a/b correlations with the appropriate red PSD.

We also found that our physically motivated HFI noise models contained a mismatch with the measured noise PSDs at the very highest frequencies (>50 Hz). Replacing the models with actual binned noise PSDs in the highest-frequency bins gave

better performance in the half-ring map test. Examples of the correlated noise spectrum and the noise model fits to the autospectra can be found in Planck Collaboration VII (2016).

To validate the results we considered the angular power spectra of both simulated and flight data half-ring map differences (i.e., splitting each pointing period into two halves, making maps from all of the first halves and second halves separately, and taking the difference of the maps), chosen because they are virtually free of all signal residuals while still supporting most of the same noise power as the full-ring data. Figure 4 illustrates the excellent agreement between the simulated and real data half-ring map *BB* spectra for 70 and 100 GHz.

3.1.3. Beams

The simulations use the so-called scanning beams (e.g., Planck Collaboration VI 2014), which give the point-spread function of for a given detector *including* all temporal data processing effects: sample integration, demodulation, ADC nonlinearity residuals, bolometric time constant residuals, etc. In the absence of significant residuals (LFI), the scanning beams may

be estimated from the optical beams by "smearing" them in the scanning direction to match the finite integration time for each instrument sample. Where there are unknown residuals in the timelines (HFI), the scanning beam must be measured directly from observations of strong point-like sources, namely planets. If the residuals are present but understood, it is possible to simulate the beam measurement and predict the scanning beam shape starting from the optical beam.

For FFP8, the scanning beams are expanded in terms of their spherical harmonic coefficients, $b_{\ell m}$, with the order of the expansion (maximum ℓ and m considered) representing a trade-off between the accuracy of the representation and the computational cost of its convolution. The LFI horns have larger beams with larger sidelobes (due to their location on the outside of the focal plane), and we treat them as full 4π beams divided into main (up to $1°9$, $1°3$, and $0°9$ for 30, 44, and 70 GHz, respectively), intermediate (up to $5°$), and sidelobe (above $5°$) components (Planck Collaboration IV 2016). This division allows us to tune the expansion orders of the three components separately. HFI horns are limited to the main beam component, measured out to $100'$ (Planck Collaboration VII 2016). Since detector beams are characterized independently, the simulations naturally include differential beam and pointing systematics.

3.1.4. Bandpasses

Both the LFI and HFI detector bandpasses are based on ground measurements (see Zonca et al. 2009; and Planck Collaboration IX 2014, respectively), although flight data processing for LFI now uses in-flight top-hat approximations rather than the ground measurements that were found to contain systematic errors. Differences in the bandpasses of detectors nominally at the same frequency (the so-called bandpass mismatch) generate spurious signals in the maps, since each detector is seeing a slightly different sky while the mapmaking algorithms assume that the signal in a pixel is the same for all detectors. To quantify the effect of these residuals, in FFP8 we generate detector timelines from foreground maps in two ways, one that incorporates the individual detector bandpasses, the other using an average bandpass for all the detectors at a given frequency.

This effect of the bandpass mismatch can be roughly measured from either flight or simulated data using so-called spurious component mapmaking, which provides noisy all-sky estimates of the observed sky differences (the spurious maps), excluding polarization, between individual detectors and the frequency average. We compare the amount of simulated bandpass mismatch to flight data in Figs. 5 (LFI) and 6 (HFI). The spurious component approach is detailed in Appendix A. Mismatch between FFP8 and flight data is driven by inaccurate bandpass description (LFI) and incomplete line emission simulation (HFI). The noisy pixels that align with the *Planck* scanning rings in the HFI maps are regions where the spurious map solution is degenerate with polarization due to insufficient observation orientations.

3.2. The foreground sky

The quality of the reconstructed CMB signal depends on the accuracy of the separation of the astrophysical foreground components from the *Planck* data, which requires building simulations as close to the complexity of the real sky as possible over a large frequency range. To this end, a complete model of multi-component sky emission, called the *Planck* sky model

(PSM), has been developed. FFP8 simulations make use of version 1.9.0 of the PSM, a major update of the published version 1.7.8 (Delabrouille et al. 2013). The foreground components generated by the PSM are diffuse Galactic emission from thermal dust, spinning dust, relativistic electrons (synchrotron emission), free-free radiation, and CO (the $J = 1 \rightarrow 0$, $J = 2 \rightarrow 1$, and $J = 3 \rightarrow 2$ lines at 115.27, 230.54, and 345.80 GHz, respectively), plus the cosmic infrared background (CIB), emission from radio sources, and the thermal and kinetic Sunyaev-Zeldovich (SZ) effects.

3.2.1. Thermal dust

The thermal dust emission is modelled using single-frequency template maps of the intensity and polarization, together with a pixel-dependent emission law.

For FFP8 the thermal dust emission templates are derived from the *Planck* 353 GHz observations. This update of the original PSM dust model, described in Delabrouille et al. (2013) and based on Finkbeiner et al. (1999), is necessary to provide a better match to the emission observed by *Planck*. While one option would be simply to use the dust opacity map obtained in Planck Collaboration XI (2014), this map still suffers from significant contamination by CIB anisotropies and infrared point sources. Using it as a 353 GHz dust template in simulations would result in an excess of small scale power (from CIB and infrared sources) scaling exactly as thermal dust across frequencies. The resulting component represents correctly neither dust alone (because of an excess of small scale power) nor the sum of dust and infrared sources (because the frequency scaling of the CIB and infrared sources is wrong). For simulation purposes, the main objective is not to have an exact map of the dust, but instead a map that has the right statistical properties. Hence we produce a template dust map at 353 GHz by removing that fraction of the small-scale power that is due to CIB emission, infra-red sources, CMB, and noise.

The emission of strong point sources is first subtracted from all *Planck* 2015 maps and from the IRAS 100 μm map following Miville-Deschênes & Lagache (2005). We mask point sources detected in *Planck* and IRAS maps with an algorithm designed specifically to avoid false detections in regions of dust emission where localized compact regions of diffuse dust emission can easily be mistaken for point sources. This algorithm detects sources against a background level computed directly on the maps in the immediate vicinity of each pixel. Specifically, for each pixel p, we compute the standard deviation $\sigma(p)$ of the observed map in a neighbourhood $N(p)$ consisting of a small ring centred on the source, between the two radii $r_1(p)$ and $r_2(p)$. We consider as part of point sources those objects for which the flux is above a threshold of 3 to 5 times $\sigma(p)$. The parameters of the filter for each frequency map – the full-width half-maximum (FWHM) radii defining the disc, and the source detection threshold – are optimized by trial and error. We then generate a mask from the union of small discs of radius 1 FWHM centred on all pixels above the detection threshold at any frequency. The gaps created by masking the point sources are filled by extrapolation of the neighbouring regions using a minimum-curvature spline gap filler. This procedure avoids most false detections of point sources in filaments of dust emission, unless there is a compact source significantly brighter than the local emission of the filament. The drawback of this approach is that extended compact sources such as nearby galaxies are not masked.

The dust intensity template is then corrected for CIB, CMB, and noise fluctuations using the multi-dimensional
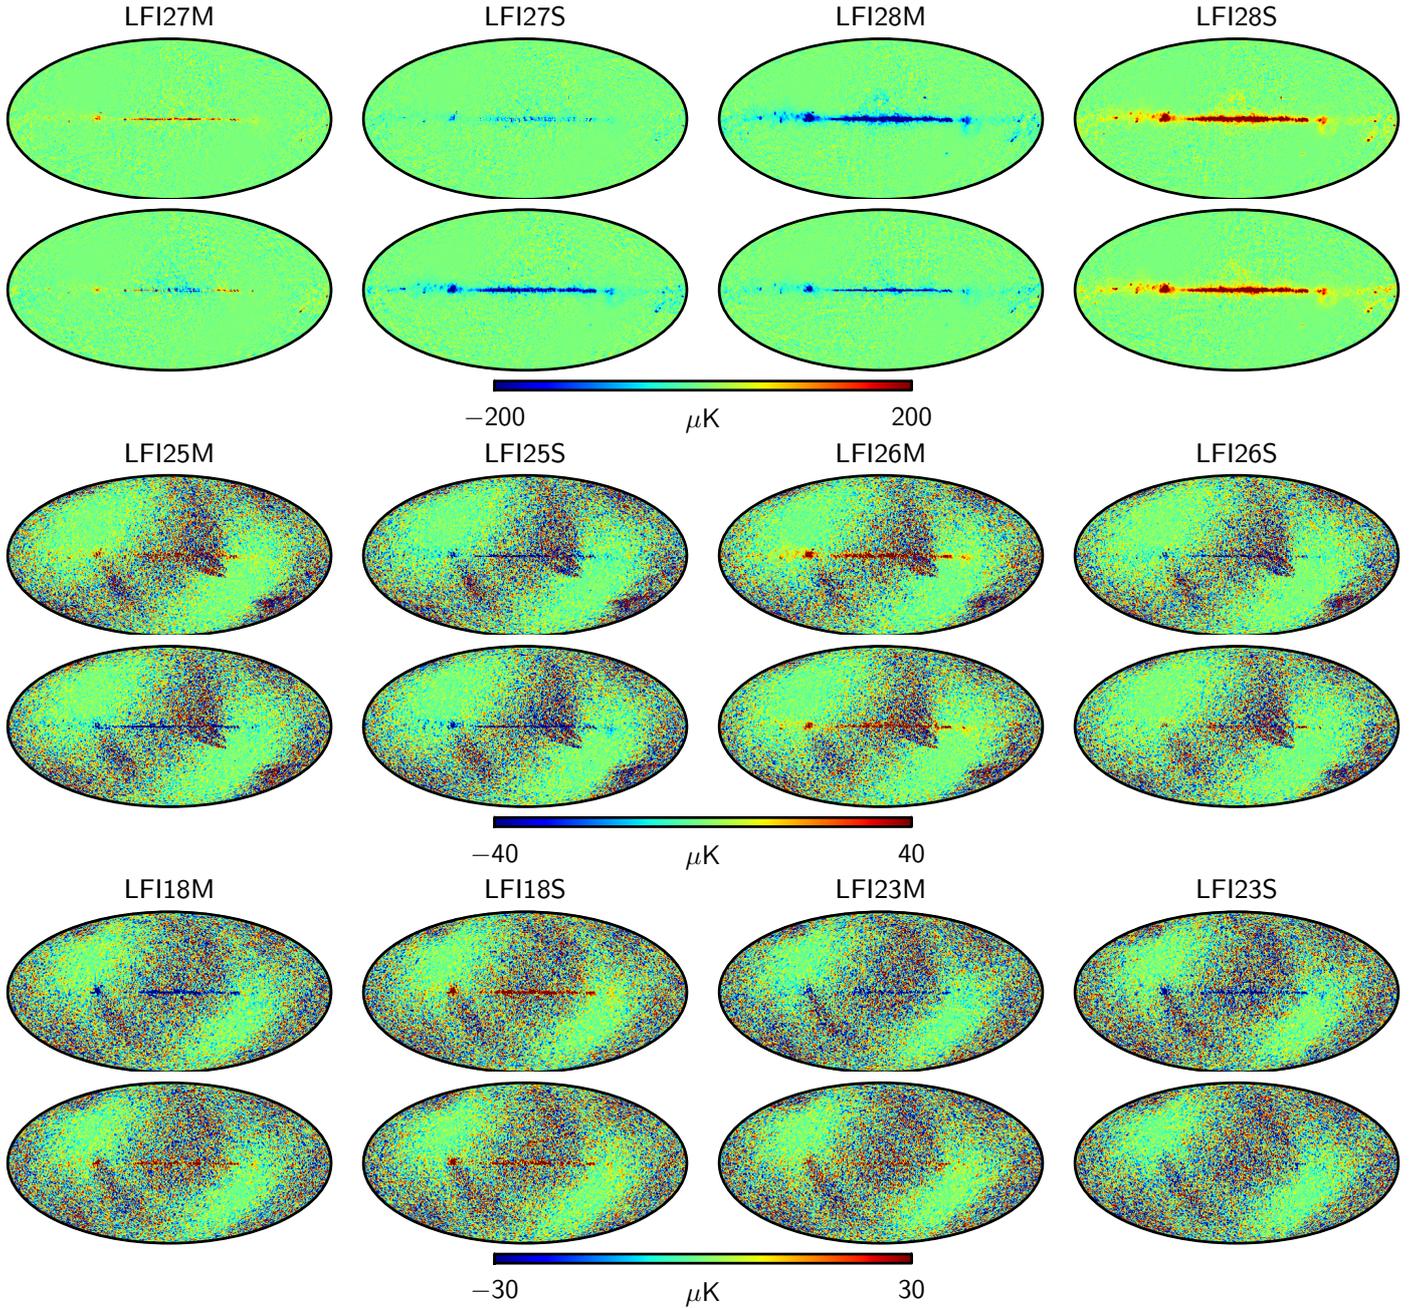
Fig. 5. Pairs of measured bandpass mismatch LFI maps, comparing FFP8 simulations (*upper*) and 2015 *Planck* (*lower*) data. Our method of estimating the total bandpass mismatch between each detector and a frequency average is detailed in Appendix A.

Generalized Needlet Internal Linear Combination (GNILC) of Remazeilles et al. (2011). Starting from the 10 de-sourced maps (nine *Planck* maps and the IRAS 100 μm map), we model the total emission at a given frequency ν as:

$$y_\nu = s_\nu + x_\nu \qquad (1)$$

where y_ν is the observed map, s_ν the map of foreground emission (mostly Galactic emission), and x_ν the sum of CMB, CIB, and noise. The total 10×10 covariance matrix R_y of the maps, in any localized region in pixel and/or harmonic space, is the sum of the foreground term R_s and a contamination term R_x due to CMB, CIB and noise,

$$R_y = R_s + R_x. \qquad (2)$$

We perform a needlet decomposition of each map as in Remazeilles et al. (2011), except that the spectral windows correspond to Gaussian bandpass filters in harmonic space. We use the best fit CMB C_ℓ from Planck Collaboration XV (2014), the CIB power spectrum at each frequency and cross-correlation coefficients from Planck Collaboration XXX (2014), and the noise level from *Planck* half-ring difference maps to obtain a model $\widehat{R_x}$ of the covariance R_x of CMB, CIB, and noise in a set of needlet domains localized on the sky. We then follow the procedure described in Remazeilles et al. (2011) to reconstruct maps of s_ν at up to 5′ resolution. For each needlet domain, the covariance of the needlet power is whitened locally as

$$\widetilde{R_y} = \widehat{R_x}^{-1/2} \widehat{R_y} \widehat{R_x}^{-1/2}. \qquad (3)$$
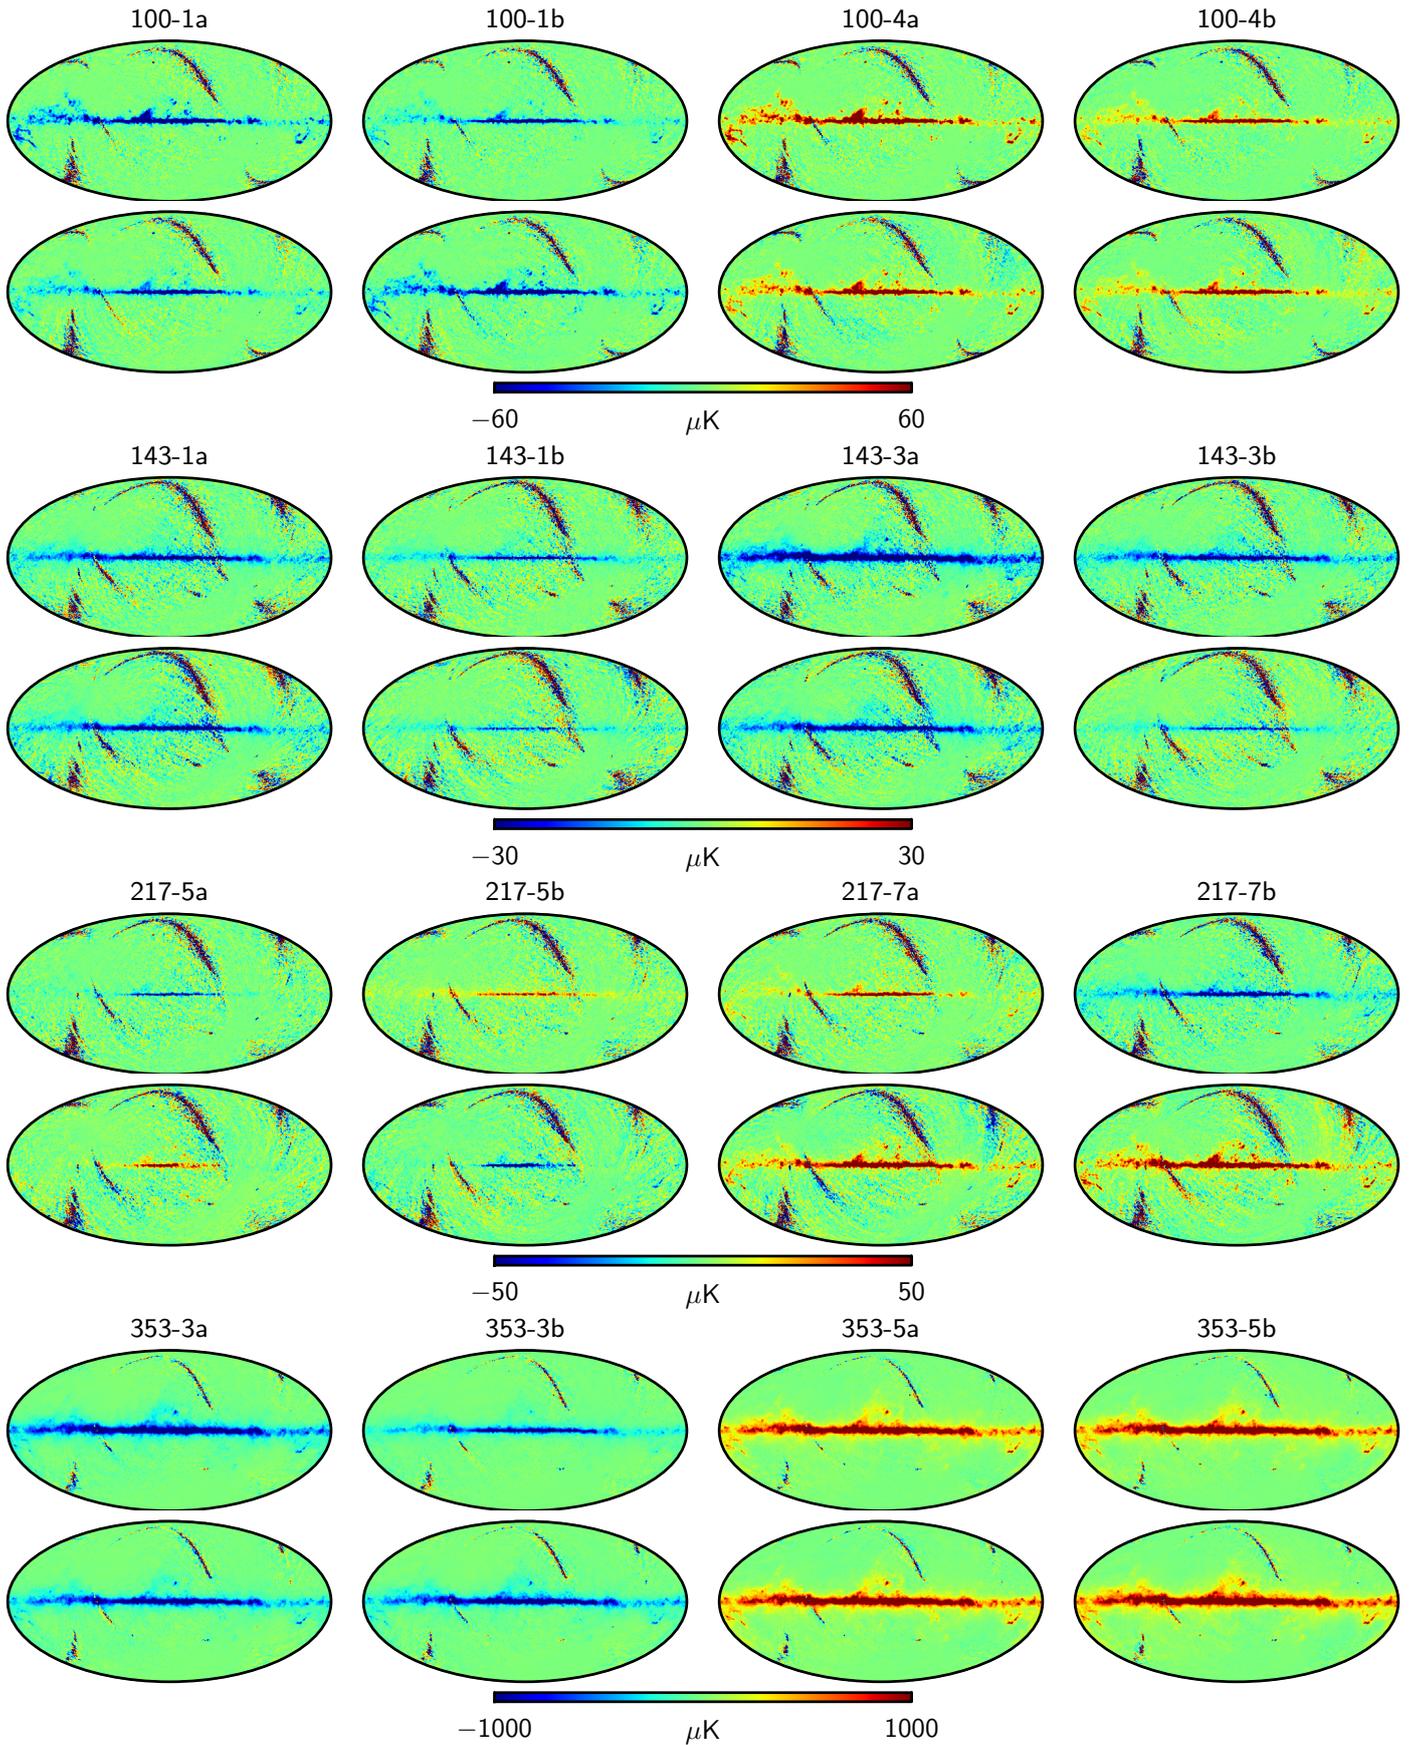
Fig. 6. Pairs of measured bandpass mismatch HFI maps, comparing FFP8 simulations (*upper*) and 2015 *Planck* (*lower*) data. Our method of estimating the total bandpass mismatch between each detector and a frequency average is detailed in Appendix A.

The matrix $\widetilde{R}_y$ is diagonalized and the Akaike information criterion (AIC; Akaike 1974) is used to filter out the CMB, CIB, and noise contamination x_ν in dimensions of the observation space where the total covariance is compatible with $s_\nu = 0$. The final map of foreground emission s_ν at each needlet scale is reconstructed by coadding the needlets for which there is significant additional emission from Galactic foregrounds according to the AIC criterion. However, the original source-subtracted 353 GHz *Planck* intensity map is used instead of the processed map in the Galactic plane, the LMC, the SMC, regions of very strong localized emission, and the extended galaxies M 31, M 33, M 81, M 82, M 87, M 101, NGC 55, NGC 300, and NGC 2403. The reason for this is that the multi-dimensional GNILC process is not perfectly local in pixel space, so that bright local structures such as the Galactic plane or nearby galaxies are analysed together with regions of significantly lower emission. This results in excessive filtering of some small-scale structure in the bright objects themselves, and insufficient filtering in the pixels located in their immediate vicinity. While the latter is not much of a problem (giving a non-crucial loss of efficiency of the multi-dimensional GNILC filtering), the former is not acceptable since it impacts the brightest objects in the sky. As CIB, CMB, and noise contributions are completely subdominant in such regions of very bright emission, it is better to use the original map than the processed one for the generation of the dust template in these regions.

Next, a mask is applied to remove from the 353 GHz dust map residuals of the positive SZ effect towards bright galaxy clusters. The mask consists of discs centred on CLASS1 (i.e., the most reliable) SZ candidates in the PSZ1 catalogue (Planck Collaboration XXIX 2014), with radius $5 \times \theta_{500}$[10]. Gaps generated in the map by this masking are filled in again by a minimum-curvature spline "inpainting". Finally, a colour-correction factor is applied to obtain a map of brightness emission at 353 GHz rather than the map integrated in the *Planck* 353 GHz channel. The 353 GHz dust map obtained in this way, which is subsequently used as a dust template, is illustrated in Fig. 7, where it is also compared with the model of Finkbeiner et al. (1999) and with the map of optical depth published in Planck Collaboration XI (2014).

The pixel-dependent emission law is derived from Planck Collaboration XI (2014), fitting *Planck* 353, 545, and 857 GHz and IRAS 100 μm data with a modified blackbody. The emission of the strongest point sources is subtracted from all maps before performing the fits, with gaps filled by the same method used for the dust intensity template.

For each pixel p, the spectral energy distribution is approximated with a modified blackbody, i.e.,

$$I_\nu(p) = I^F_{353\,\text{GHz}}(p)\left(\frac{\nu}{353\,\text{GHz}}\right)^{\beta_{\text{FIR}}(p)}\left(\frac{B_\nu(T_{\text{dust}}(p))}{B_{353\,\text{GHz}}(T_{\text{dust}}(p))}\right), \qquad (4)$$

where the maps of the dust temperature T_{dust} and spectral index β_{FIR} are those presented in Planck Collaboration XI (2014). The sky map $I^F_{353\,\text{GHz}}$ is produced by applying the multi-dimensional GNILC of Remazeilles et al. (2011).

Maps of the Q/I and U/I ratios between Stokes parameters were computed from the best versions of the *Planck* maps available at the time of generation of the simulations in June 2014, smoothed to 30′ resolution to reduce noise. At 353 GHz, the observed maps are used for polarization, while the intensity map is obtained from multi-dimensional GNILC processing. All maps

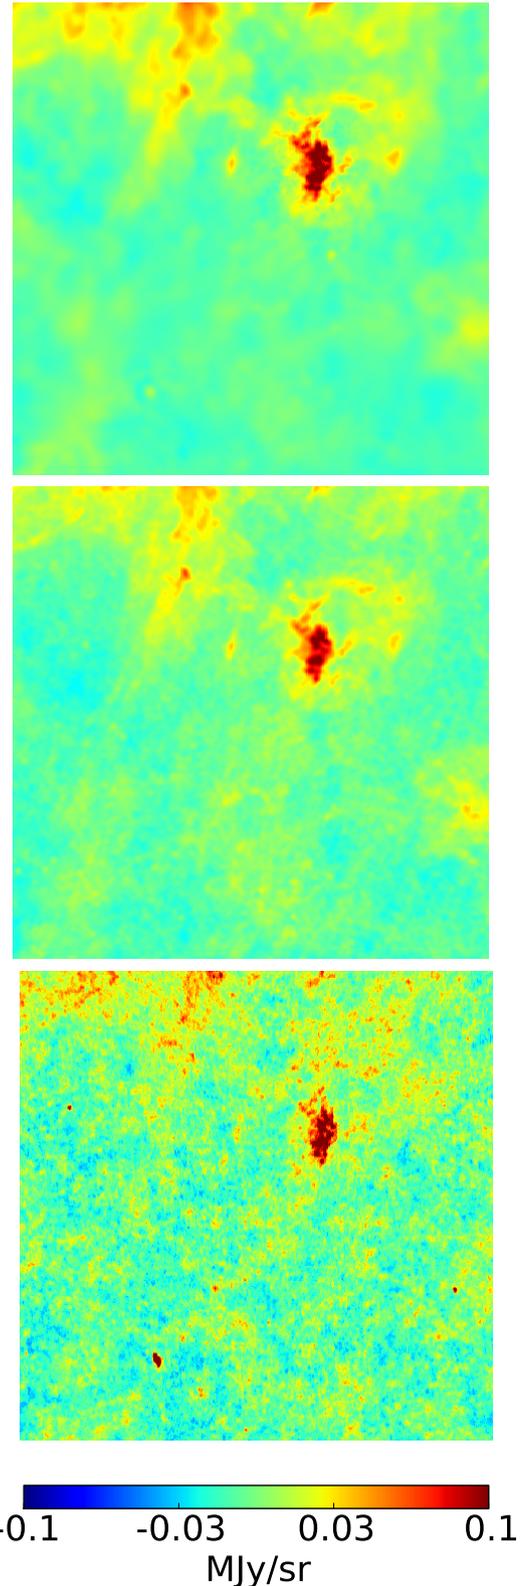[10] θ_{500} is the observed radius $R_{500}/D_A(z)$, where $D_A(z)$ is the angular diameter distance of the cluster.

Fig. 7. $12\overset{\circ}{.}5 \times 12\overset{\circ}{.}5$ gnomonic projection of dust intensity templates centred at high latitude $(l, b) = (90°, -80°)$ at a map resolution of 5′. *Top*: new PSM dust from multi-dimensional GNILC (Remazeilles et al. 2011). *Middle*: original PSM dust (Delabrouille et al. 2013) based on the model of Finkbeiner et al. (1999). *Bottom*: *Planck* 2013 dust model (Planck Collaboration XI 2014). Comparing the top and bottom panels, the improvement in reducing the contamination of the dust by small-scale CIB anisotropies at high Galactic latitude is clearly visible.
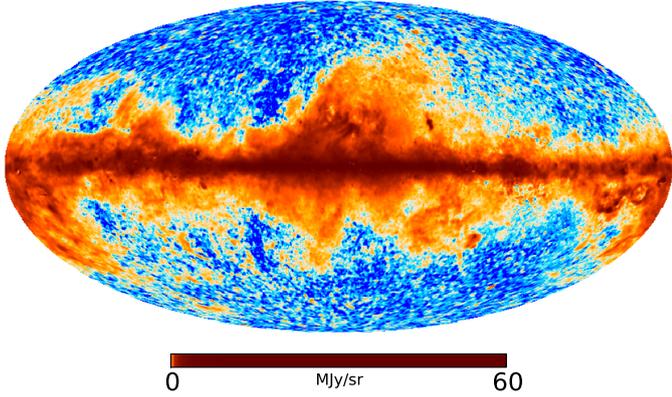
Fig. 8. Template of spinning dust intensity.

are corrected for the dust bandpass mismatch leakage. Structure on small scales is added in order to have no break in the power spectra of dust polarization. These maps determine the polarization orientation and fraction. For each frequency, the Stokes Q and U maps of the model are then obtained by multiplying the Q/I and U/I maps by the total intensity map at full resolution. The polarization orientation does not depend on frequency, but the difference between β_{mm} and β_{FIR} introduces a spectral dependence of the polarization fraction, which varies over the sky.

To extrapolate the model to frequencies below 353 GHz we follow Planck Collaboration Int. XVII (2014), using a spectral index map β_{mm} slightly smaller than β_{FIR} for the total dust intensity, but not for the polarized emission. The difference between β_{mm} and β_{FIR} is a quadratic function of T_{dust}, which has been fitted comparing the *Planck* maps at 143 and 353 GHz, after subtraction of the CMB anisotropies. It introduces a small spectral difference between dust intensity and polarization that varies over the sky, similar to that reported in Planck Collaboration Int. XXII (2015).

3.2.2. Spinning dust

Several observations indicate an excess of microwave emission in the 10−100 GHz range (Kogut et al. 1996; Leitch et al. 1997; de Oliveira-Costa et al. 1999; Watson et al. 2005). This emission has been proposed to be the dipole radiation from rapidly spinning nano-scale dust grains (Draine & Lazarian 1998). Detailed analyses of the data from various experiments show detections of a component that is consistent with the spinning dust model (Davies et al. 2006; Miville-Deschênes et al. 2008; Kogut et al. 2011; Planck Collaboration XX 2011; Planck Collaboration Int. XV 2014; Planck Collaboration X 2016).

The spinning dust map used for FFP8 simulations is a simple realization of the spinning dust model, post-processed to remove negative values occurring in a few pixels because of the generation of small-scale fluctuations on top of the spinning dust template extracted from WMAP data by Miville-Deschênes et al. (2007). The model uses a single template map, illustrated in Fig. 8, which is scaled by a single emission law parameterized following the model of Draine & Lazarian (1998).

3.2.3. Synchrotron

At low frequencies (below about 80 GHz) the dominant contaminant of the polarized CMB signal is the Galactic synchrotron emission that arises from the acceleration of relativistic cosmic rays in the Galactic magnetic field. FFP8 synchrotron emission

is modelled on the basis of the template emission map observed at 408 MHz by Haslam et al. (1982). This template synchrotron map is extrapolated in frequency using a spectral index map corresponding to a simple power law.

3.2.4. Free-free

Electron-ion interactions in the ionized phase of the ISM produce emission that is in general fainter than both the synchrotron and the thermal dust emission outside of the active star-forming regions in the Galactic plane. The free-free model uses a single template, which is scaled in frequency by a specific emission law. The free-free spectral index is a slowly varying function of frequency and depends only slightly on the local value of the electron temperature T_e (Bennett et al. 1992, 2003; Dickinson et al. 2003). In FFP8, the free-free spectral dependence is modelled by assuming a constant electron temperature $T_e = 7000$ K, as described in Dickinson et al. (2003).

3.2.5. Cosmic infrared background

The CIB model relies on the distribution of individual galaxies in template maps based on the distribution of dark matter at a range of relevant redshifts. We assume the CIB galaxies can be grouped into three different populations (proto-spheroid, spiral, starburst) as described in Cai et al. (2013). Within each population, galaxies have the same SED, while the flux density is randomly distributed according to redshift-dependent number counts obtained from JCMT/SCUBA-2 observations (Chen et al. 2013) and the *Planck* ERCSC (Negrello et al. 2013), as well as observations cited in Cai et al. (2013), most notably *Herschel*/Spire (Hermes; Béthermin et al. 2013) and AzTEC/ASTE (Combined ACES; Scott et al. 2012).

We use the Class (Blas et al. 2011) software to generate dark matter maps at 17 different redshifts between 1 and 5.5. Since the galaxy distribution does not exactly follow the dark matter distribution, we modify the a_{lm}^{DM} coefficients of dark matter anisotropies given by Class at redshift z_i such that

$$a_{\ell m,j}(z_i) = a_{\ell m}^{DM}(z_i) \times b_j(z_i) \times \left(1 + \left(\frac{\ell}{\beta}\right)^{\alpha}\right), \tag{5}$$

where $b_j(z)$ is the population-dependent component of the bias of the population j, and is identical for spirals and starbursts. $\left(1 + \left(\frac{\ell}{\beta}\right)^{\alpha}\right)$ is the scale-dependent part of the bias. Values of $\alpha = 1.44$ and $\beta = 2654$ were obtained by fitting *Planck* spectra. Template maps generated from the $a_{\ell m,j}(z_i)$ coefficients are then exponentiated to avoid negative pixels.

Galaxies are randomly distributed with a probability of presence proportional to the pixel values of the template maps. One map is generated for each population, at each redshift, and associated with a redshifted SED depending on the population. The emission of these maps (initially at a reference frequency) can be extrapolated to any frequency using the associated redshifted SED. By summing the emission of all maps, we can generate CIB maps at any frequency in the range of validity of our model.

A comparison of CIB power spectra of the simulated maps and the maps obtained by *Planck* observations at HFI frequencies is shown in Fig. 9; agreement is excellent at all but the lowest two HFI frequencies where the CIB is subdominant. See Planck Collaboration XXX (2014) for further discussion.
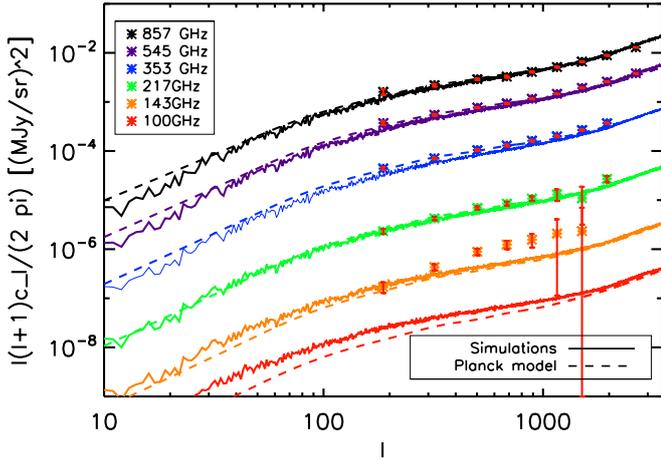
Fig. 9. Simulated and observed cosmic infrared background power spectra. Solid lines are from the PSM simulations; dashed lines are the *Planck* "extended halo model" which includes the correlated and Poisson contributions to the CIB power spectrum from dusty galaxies. Data points are *Planck* observations, corrected for the contribution of radio sources, all from Planck Collaboration XXX (2014).

3.2.6. Radio sources

Radio source modelling in the PSM was described in detail in Delabrouille et al. (2002). Here, we describe only the main changes since that pre-launch version.

For "strong" radio sources ($S_{30} > 0.5$ Jy), we use radio sources compiled by González-Nuevo et al. (2008) at 0.84, 1.4, or 4.85 GHz. For sources observed at two of these frequencies, we extrapolate or interpolate to the third frequency assuming the spectral index estimated from two observed. For sources observed at only one frequency, we use differential source counts from de Zotti et al. (2010) to obtain the ratio of steep- to flat-spectrum sources in each interval of flux density considered. From this ratio, we assign spectral indices (randomly) to each source within each flux density interval. Fiducial Gaussian spectral index distributions as a function of spectral class are obtained from the literature (Massardi et al. 2011; Sadler et al. 2008). These are then adjusted slightly until there is reasonable agreement between the PSM differential counts and the model counts predicted by de Zotti et al. (2010). The final means and standard deviations of these spectral index distributions are -0.7 ± 0.2 for steep-spectrum sources, -0.3 ± 0.2 for flat-spectrum radio quasars (FSRQ), and -0.1 ± 0.2 for BL Lacertae objects (BL Lacs).

For "faint" radio sources ($S_{30} \leq 0.5$ Jy), the pre-launch PSM showed a deficit of sources resulting from inhomogeneities in surveys at different depths. We address this issue by constructing a simulated catalogue of sources from the model differential counts of de Zotti et al. (2010) at 1.4 GHz. We replace the simulated sources by the observed ones, wherever possible. If, however, in any particular pixel, we have a shortfall of observed sources, we make up the deficit with the simulated sources. Every source in this new catalogue is given a model-derived spectral class. We thus assign a spectral index to each source based on the spectral class, and model the spectrum of each source using four power laws covering the ranges $\nu < 5$ GHz, $5 < \nu < 20$ GHz, $20 < \nu < 100$ GHz, and $\nu > 100$ GHz. We also assume some steepening of the spectral index with frequency, with fiducial values of the steepening obtained from the literature (Massardi et al. 2011; Sadler et al. 2008). The steepening
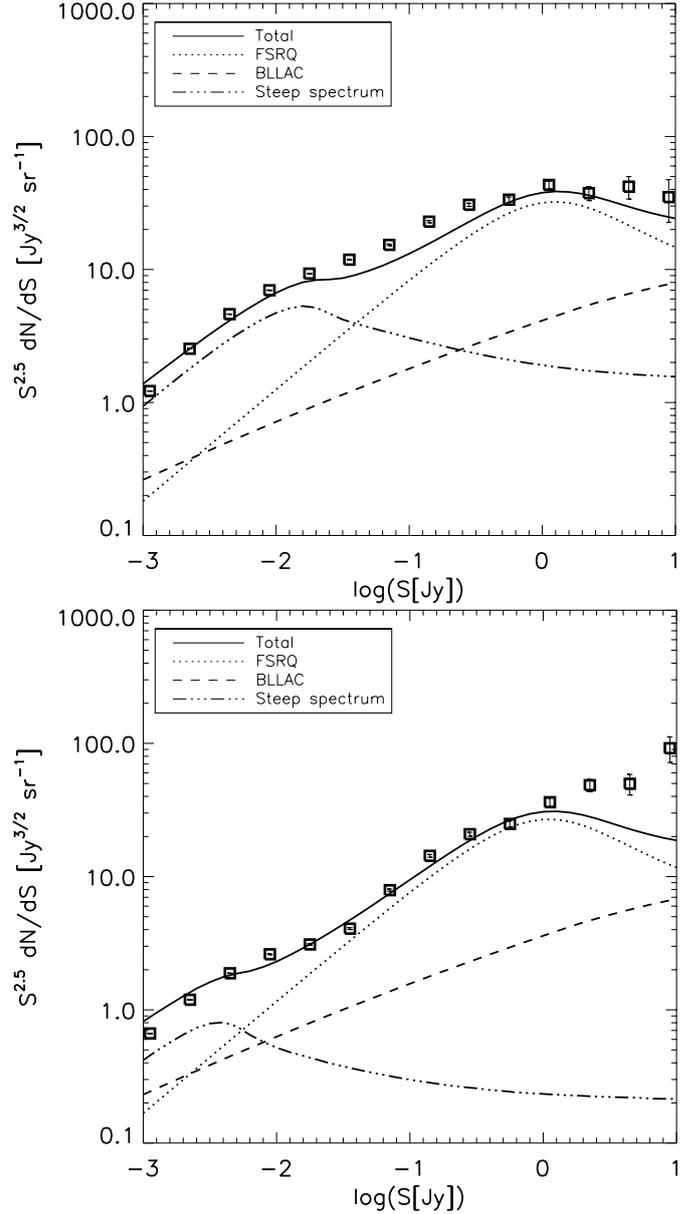

Fig. 10. Counts of the faint and the strong radio sources at 30 GHz (*top*) and 100 GHz (*bottom*). The differential counts of the faint and the strong radio sources between 0.005 Jy and 1 Jy are shown by vertical bars.

of the spectral index assumed for 5–20 GHz is -0.3 for steep-spectrum sources and -0.3 for flat-spectrum sources (FRSQs and BL Lacs). For 20–100 GHz, we assume no further steepening for steep-spectrum sources, and steepening of -0.3 for FRSQs and BL Lacs.

We combine the faint and strong radio source catalogues we constructed and compute the differential source counts on these sources between 0.005 Jy and 1 Jy (see Fig. 10). These show good agreement with the model differential counts predicted by de Zotti et al. (2010), particularly at the faintest end. Finally we also model the polarization of these radio sources using the measured polarization fractions from Ricci et al. (2004), comprising 71 steep- and 126 flat-spectrum sources; for each simulated source we draw a polarization fraction at random from the list of real sources of the same spectral type.

Table 3. Cosmological parameter values from the PR1-2013 and PR2-2015 data releases, used in FFP8 and FFP8.1, respectively.

Parameter	Symbol	PR1-2013 (FFP8)	PR2-2015 (FFP8.1)
Baryon density	$\omega_b = \Omega_b h^2$	0.0222	0.0223
Cold dark matter density	$\omega_c = \Omega_c h^2$	0.1203	0.1184
Neutrino energy density	$\omega_\nu = \Omega_\nu h^2$	0.00064	0.00065
Dark energy density	Ω_Λ	0.6823	0.6931
Hubble parameter, $H_0 = 100\, h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ . . .	h	0.6712	0.6787
Primordial curvature perturbation spectrum			
amplitude (at $k = 0.05\,\mathrm{Mpc^{-1}}$)	A_s	2.09×10^{-9}	2.14×10^{-9}
spectral index	n_s	0.96	0.97
Thomson optical depth through reionization	τ	0.065	0.067

3.2.7. SZ clusters

SZ clusters are simulated following the model of Delabrouille, Melin, and Bartlett (DMB) as implemented in the PSM (Delabrouille et al. 2002). A catalogue of halos is drawn from a Poisson distribution of the mass function from Tinker et al. (2008) with a limiting mass of $M_{500,\mathrm{true}} > 2 \times 10^{13}\, M_\odot$ [11]. We use the pressure profile from Arnaud et al. (2010) to model the thermal SZ emission of each halo given its redshift and mass, $M_{500,x}$. In order to match the observed counts in *Planck*, we set $M_{500,x}/M_{500,\mathrm{true}} = 0.65$. Fixing this ratio to unity would lead to twice as many clusters simulated as appear in the actual data, as shown in Planck Collaboration XX (2014). We determine the cluster temperature from the $M_{500,x}-T$ relation of Arnaud et al. (2005), and assume that the profiles are isothermal. These steps allow us to compute the first-order thermal relativistic correction (Itoh et al. 1998) and the kinetic SZ effect for each cluster, both of which are included in the simulation. Finally, we inject MCXC (Piffaretti et al. 2011) and MaxBCG (Koester et al. 2007) clusters following the same model, and remove from the simulation corresponding clusters in each redshift and mass range. Hence the SZ simulation features the majority of known X-ray and optical clusters, and is fully consistent with X-ray scaling laws and observed *Planck* SZ counts.

3.2.8. Total foreground sky

The sky model is simulated at a resolution common to all components by smoothing the maps with an ideal Gaussian beam of FWHM of $4'$. The `HEALPix`[12] pixelization in Galactic coordinates is used for all components, with $N_{\mathrm{side}} = 2048$ and $\ell_{\mathrm{max}} = 6000$. Sky emission maps are generated by numerically band-integrating the sky model maps (emission law of each component, in each pixel) over the frequency bands both of each detector in the focal plane and – using an average over the detectors at a given frequency – of each channel. The band-integrated maps are essentially observations of the model sky simulated by an ideal noiseless instrument with ideal Gaussian beams of FWHM equal to the resolution of the model sky.

3.3. The CMB sky

The CMB sky is simulated in three distinct components, namely lensed scalar, tensor, and non-Gaussian complement. The total

---

[11] M_{500} is the total cluster mass within a sphere of radius R_{500}, defined as the radius within which the mean mass over-density of the cluster is 500 times the cosmic critical density at the redshift of the cluster. $M_{500} = (4\pi/3)R_{500}3[500\rho_c(z)]$, with $\rho_c(z) = 3H^2(z)/(8\pi G)$, where $H(z)$ is the Hubble parameter with present-day value H_0.
[12] http://healpix.sourceforge.net

CMB sky is then the weighted sum with weights 1, $\sqrt{r}$, and f_{NL}, respectively. For FFP8, all CMB sky components are produced as spherical harmonic representations of the I, Q, and U skies.

The FFP8 CMB sky is derived from our best estimate of the cosmological parameters available at the time of its generation, namely those from the first *Planck* data release (PR1-2013; Planck Collaboration I 2014), augmented with a judicious choice of reionization parameter τ, as listed in Table 3.

3.3.1. The scalar CMB sky

The scalar component of the CMB sky is generated including lensing, Rayleigh scattering, and Doppler boosting effects. Using the `camb` code (Lewis et al. 2000), we first calculate fiducial unlensed CMB power spectra C_ℓ^{TT}, C_ℓ^{EE}, C_ℓ^{TE}, the lensing potential power spectrum $C_\ell^{\phi\phi}$, and the cross-correlations $C_\ell^{T\phi}$ and $C_\ell^{E\phi}$. We then generate Gaussian T, E, and ϕ multipoles with the appropriate covariances and cross-correlations using a Cholesky decomposition and three streams of random Gaussian phases. These fields are simulated up to $\ell_{\mathrm{max}} = 5120$. In addition, we add a dipole component to ϕ to account for the Doppler aberration due to our motion with respect to the CMB (Challinor & van Leeuwen 2002):

$$\phi_\beta(\hat{n}) = \boldsymbol{\beta} \cdot \hat{n}, \tag{6}$$

where $\boldsymbol{\beta}$ is in the direction of Galactic coordinates $(l, b) = (264\overset{\circ}{.}4, 48\overset{\circ}{.}4)$ with amplitude $|\boldsymbol{\beta}| = 0.00123$, corresponding to a velocity of $369\,\mathrm{km\,s^{-1}}$ (Hinshaw et al. 2009).

We compute the effect of gravitational lensing on the temperature and polarization fields, using an algorithm similar to `LensPix` (Lewis 2005). We use a fast spherical harmonic transform to compute the temperature, polarization, and deflection fields:

$$[d_\theta + \mathrm{i}d_\varphi] = \sum_{\ell m} -\sqrt{\ell(\ell + 1)}\, {}_1Y_{\ell m}(\hat{n})\, \phi_{\ell m}. \tag{7}$$

The unlensed CMB fields T, Q, and U are evaluated on an equicylindrical pixelization (ECP) grid with $N_\theta = 32\,768$ and $N_\varphi = 65\,536$, while the deflection field is evaluated on a `HEALPix` $N_{\mathrm{side}} = 2048$ grid. We then calculate the "lensed positions" $\hat{n}'$ for each $N_{\mathrm{side}} = 2048$ `HEALPix` pixel. In the flat-sky limit these are given schematically as $\hat{n}' = \hat{n} + d(\hat{n})$; in practice on the curved sky we parallel transport $\hat{n}$ a distance $|d(\hat{n})|$ along the great circle in the direction $d(\hat{n})$. We then interpolate T, Q, U at the lensed positions using 2D cubic Lagrange interpolation on the ECP grid.

After lensing, we incorporate the frequency-dependent Doppler modulation effect (Planck Collaboration XXVII 2014), multiplying the temperature and polarization maps by

Table 4. Boost factors.

Frequency [GHz]	Boost factor b_ν
30	1.05
44	1.1
70	1.25
100	1.51
143	1.96
217	3.07
353	5.38
545	8.82
857	14.20

$[1 + b_\nu \boldsymbol{\beta} \cdot \hat{\boldsymbol{n}}]$, with b_ν factors given in Table 4. These boost factors incorporate bandpass corrections calculated given the frequency-dependence of the Doppler effect. We then evaluate lensed, Doppler boosted $T_{\ell m}$, $E_{\ell m}$, and $B_{\ell m}$ up to $\ell_{\max} = 4\,096$ with a harmonic transform of the $N_{\text{side}} = 2048$ HEALPix map of these interpolated T, Q, and U values.

We next add frequency-dependent Rayleigh scattering effects, using the "R" matrices defined in Eq. (7) of Lewis (2013). We use effective central frequencies for the ν^4 terms of 30, 44, 70, 105, 148, 229, 372, 577, and 891 GHz for the 30–857 GHz channels, and effective central frequencies for the ν^6 terms of 30, 44, 70, 108, 151, 233, 378, 585, and 905 GHz.

Finally we add a second-order temperature quadrupole following the equations in Kamionkowski & Knox (2003). Since the main *Planck* data processing removes the frequency-independent part (Planck Collaboration VIII 2016), we simulate only the residual frequency-dependent temperature quadrupole. After subtracting the frequency-independent part, the simulated quadrupole has frequency dependence $\propto (b_\nu - 1)/2$, which we calculate using the bandpass-integrated b_ν boost factors given in Table 4. Defining the z-axis to be in the direction of the solar dipole, at each frequency we add

$$a_{20}^T = \frac{4(b_\nu - 1)}{6} \sqrt{\frac{\pi}{5}} \, \beta^2 \, T_0. \tag{8}$$

3.3.2. The tensor CMB sky

In addition to the scalar CMB simulations, we also generate a set of CMB skies containing primordial tensor modes. Using the fiducial cosmological parameters of Table 3, we calculate the tensor power spectra $C_\ell^{TT,\text{tensor}}$, $C_\ell^{EE,\text{tensor}}$, and $C_\ell^{BB,\text{tensor}}$ using camb with a primordial tensor-to-scalar power ratio $r = 0.2$ at the pivot scale $k = 0.05\,\text{Mpc}^{-1}$. We then simulate Gaussian T, E, and B-modes with these power spectra, and convert these to spherical harmonic representations of the corresponding I, Q and U maps. Note that the default $r = 0.2$ means that building the FFP8a-d maps requires rescaling each CMB tensor map by $\sqrt{r/0.2}$ for each of the values of r in Table 2.

3.3.3. The non-Gaussian CMB sky

We use a new algorithm to generate simulations of CMB temperature and polarization maps containing primordial non-Gaussianity. Non-Gaussian fields in general have a non-vanishing bispectrum contribution sourced by mode correlations. The bispectrum, the Fourier transform of the 3-point correlation function, can then be characterized as a function of three wavevectors, $F(k_1, k_2, k_3)$. Depending on the physical

mechanism responsible for generating the non-Gaussian signal, it is possible to introduce broad classes of model that are categorized by the dependence of F on the type of triangle formed by the three momenta k_i (see e.g., Babich et al. 2004). Here, we focus on non-Gaussianity of local type, where the bulk of the signal comes from squeezed triangle configurations, $k_1 \ll k_2 \approx k_3$. This is typically predicted by multi-field inflationary models (e.g., Moroi & Takahashi 2001; Enqvist & Sloth 2002; Lyth et al. 2003). In this scenario, the primordial gravitational potential Φ is defined in terms of a Gaussian auxiliary field Φ_{L},

$$\Phi = \Phi_{\text{L}} + f_{\text{NL}} \left(\Phi_{\text{L}}^2 - \langle \Phi_{\text{L}} \rangle^2 \right), \tag{9}$$

where higher-order terms are neglected (Verde et al. 2000; Wang & Kamionkowski 2000). Here, the non-Gaussian contribution to the primordial potential is parameterized by means of the scalar prefactor f_{NL}. The spherical harmonic coefficients of the CMB anisotropies are then given by a linear equation, the line-of-sight integral (Liguori et al. 2003)

$$a_{\ell m}^X = \frac{2}{\pi} \int \mathrm{d}r \, r^2 \Phi_{\ell m}(r) \alpha_\ell^X(r)$$
$$= \mathsf{M}_\ell^X \Phi_{\ell m}, \tag{10}$$

where we introduce the real space transfer function $\alpha_\ell^X(r)$ for temperature $(X = T)$ and polarization $(X = E)$. We also make use of a compact matrix notation for the integral that, for all practical purposes, is evaluated numerically on a finite number of grid points.

Our aim is to augment the Gaussian CMB simulations described above with non-Gaussian templates, defined by the second term in Eq. (9). Given the noiseless Gaussian map, we first compute a Wiener filter reconstruction of the primordial potential Φ_{L} on the full sky. Let P be the covariance matrix of Φ_{L} on the radial grid (e.g., Liguori et al. 2003),

$$\mathsf{P} = \left\langle \Phi_{\text{L}\,\ell_1 m_1}(r_1) \, \Phi_{\text{L}\,\ell_2 m_2}^*(r_2) \right\rangle$$
$$= 4\pi \, \delta_{\ell_2}^{\ell_1} \, \delta_{m_2}^{m_1} \int \mathrm{d}k \, \frac{\Delta_\Phi^2(k)}{k} \, j_{\ell_1}(kr_1) \, j_{\ell_2}(kr_2), \tag{11}$$

where $\Delta_\Phi^2(k)$ is the primordial power spectrum and the j_ℓ are spherical Bessel functions. Further denoting C as the CMB covariance matrix given by the fiducial power spectrum of the Gaussian simulated map s, then

$$\Phi_{\text{L}}^{\text{WF}}(s) = \mathsf{P} \mathsf{M}^\dagger \mathsf{C}^{-1} s. \tag{12}$$

Since the Wiener filter is a biased estimate, we now supplement it with a fluctuation term to yield a Gaussian constrained realization (Bertschinger 1987). The covariance of the fluctuation term is completely determined by the kernel of M,

$$\mathsf{P}^{\text{fluc}} = \mathsf{P} - \mathsf{P} \mathsf{M}^\dagger \mathsf{C}^{-1} \mathsf{M} \mathsf{P}. \tag{13}$$

We finally obtain an unbiased reconstruction of the linear gravitational potential for each simulation s,

$$\Phi_{\text{L}}(s) = \Phi_{\text{L}}^{\text{WF}}(s) + \left(\mathsf{P}^{\text{fluc}} \right)^{1/2} g, \tag{14}$$

where g is a vector of univariate Gaussian random numbers. With the reconstruction in hand, it is now straightforward to obtain the non-Gaussian template s_{NG},

$$s_{\text{NG}} = \mathsf{M} \left(\Phi_{\text{L}}(s)^2 - \langle \Phi_{\text{L}} \rangle^2 \right). \tag{15}$$

We evaluate all line-of-sight integrals (Eq. (10)) on 70 radial shells, where quadrature weights and node positions are optimized as described in Elsner & Wandelt (2009). The resulting non-Gaussian map is exact in the sense that all higher-order correlation functions are correctly represented.

3.3.4. The FFP8.1 CMB skies

The FFP8 simulations are an integral part of the analyses used to derive PR2-2015, and so were necessarily generated prior to determining that release's cosmological parameters. As such there is inevitably a mismatch between the FFP8 and the PR2-2015 cosmologies, which we address in two ways. The quick-and-dirty fix is to determine a single rescaling factor that minimizes the difference between the PR1-2013 and PR2-2015 TT power spectra and apply it to all of the FFP8 CMB maps; this number is determined to be 1.0134, and the rescaled maps have been used in several repeat analyses to confirm the robustness of various PR2-2015 results.

More rigorously though, we also generate a second set of CMB realizations based on the PR2-2015 cosmology, dubbed FFP8.1, and perform our reanalyses using these in place of the FFP8 CMB skies in both the fiducial and MC realizations. Table 3 lists the cosmological parameters used for FFP8.1 while Table 1 enumerates the current status of the FFP8.1 CMB MCs.

4. FFP8 pipelines

The two *Planck* DPCs have distinct internal data formats and access protocols, including elements that are either impossible or impractical to port to supercomputing facilities. Moreover, data access can be a significant computational bottleneck, especially in an HPC environment, so aggressive IO optimization is necessary for the most compute-intensive operations. To address these issues, we use the *Planck* exchange format[13] (EF), developed for inter-DPC exchanges and public releases, for all time-domain data, adding the necessary interfaces to EF files to the TOAST data abstraction layer. TOAST allows any analysis code to access any supported data set – including *Planck* – without worrying about the details of the underlying file formats or data distributions on disk. Instead the path to, format of, and metadata about, each data file required in a particular analysis is stored in an XML document called a run configuration file, or "runconfig". Any request by a TOAST-interfaced analysis code for data is then passed to TOAST to resolve and satisfy, using the information in the runconfig. This data abstraction also enables TOAST to cache heavily re-used data (such as the telescope pointing information) and to generate synthetic data on-the-fly (bypassing the crippling IO costs associated with performing distinct time-domain simulation and analysis phases), and to do both of those things transparently to the calling analysis code. Since the general analysis of large, correlated data sets involves global operations on all of the data, requiring very massive parallelism, TOAST uses a hybrid MPI/OpenMP programming model to enable these operations to be efficiently executed on the architectures of the current largest-scale computing resources. All of the codes used in the various FFP8 pipelines are interfaced to TOAST.

4.1. Fiducial realizations

The FFP8 fiducial realization is generated in two steps: (1) simulation of the full mission TOD for every detector; and (2) calculation of maps from the various detector subsets, intervals, and data cuts (Fig. 11). Simulation of explicit TODs allows us to incorporate each detector's full beam (including its far sidelobes)

---

[13] The file format used in DPC exchanges is similar to the official released timelines (Planck Collaboration 2015) except that detector pointing is not supplied but rather interpolated on-the-fly from 8 Hz satellite attitude files.
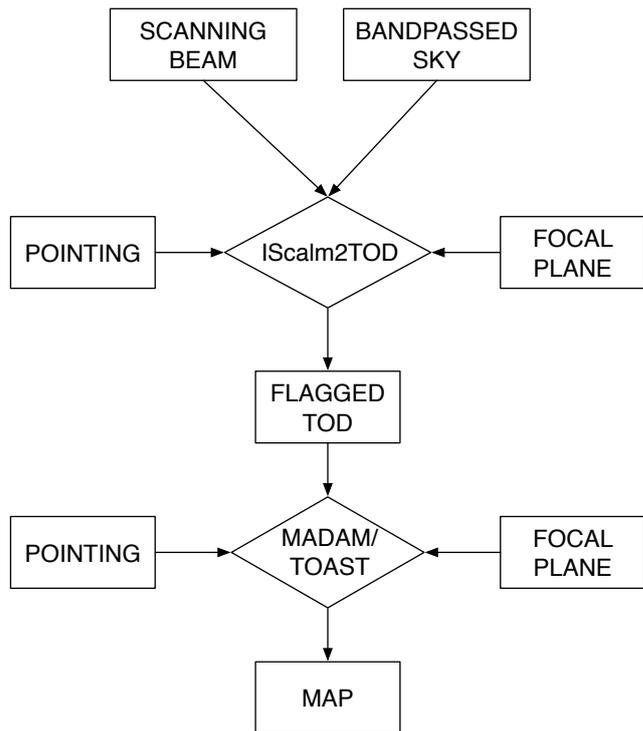

Fig. 11. Fiducial pipeline schematic.

and unique input sky (including its bandpass). As noted above, the fiducial realization is generated in six separate components – the three CMB components (lensed scalar, tensor, and non-Gaussian complement), two foreground realizations (with and without bandpass mismatch), and noise. The first five of these are simulated as explicit TODs and then mapped, while the noise is generated using the on-the-fly approach described in the noise MC subsection below.

TOD generation for any detector proceeds by first convolving the appropriate sky component with the beam at every point in a uniformly sampled data cube of Euler angle triplets (encoding the pointing and polarization orientation) to produce the "beamskyset", and then generating the time-ordered data by interpolating over the beamskyset data cube to the exact pointing and polarization orientation of each sample. Previous FFP simulations, including FFP6, accompanying the 2013 *Planck* data release, used the LevelS software package (Reinecke et al. 2006) to do this, using the conviqt (Prezeau & Reinecke 2010) and multimod tools, respectively. However, this required format conversions for the input pointing data (from exchange to LevelS) and the output time-ordered data (from LevelS to exchange), at significant IO and disk space costs. For FFP8 we have therefore embedded the critical parts of these routines into a new code, IScalm2TOD, which uses TOAST to interface directly with exchange format data. As an additional benefit, by combining the two steps into a single code the beamskyset never has to be written to disk, further reducing the IO and disk space costs. This approach also allows us trivially to handle the LFI multi-component beams (Sect. 3.1.3), where each beam is expressed at a different (ℓ_{max}, m_{max}) resolution. Previously this meant generating distinct TODs and adding them only after they had been written, tripling the IO and disk space requirements, whereas now they are generated separately but combined within IScalm2TOD prior to being written.

All of the FFP8 fiducial maps are produced using `Madam/TOAST`, a `TOAST` port of the `Madam` generalized destriping code (Keihänen et al. 2005, 2010), which allows for destriping with an arbitrary baseline length, with or without a prior on the baseline distribution (or noise filter). `Madam` is used to produce the official LFI maps, and its destriping parameters can be chosen so that it reproduces the behaviour of `Polkapix`, the official HFI mapmaking code. Comparison of the official maps and `Madam/TOAST` maps run using exchange data show that mapmaker differences are negligible compared to small differences in pointing and (for HFI) dipole subtraction that do not impact the simulation. The sky components are mapped from the TODs generated by `IScalm2TOD`, while the fiducial noise is taken to be realization 10 000 of the noise MC (with realizations 0000−9999 reserved for the noise MC itself). Summarizing the key differences in the map making parameters for each *Planck* frequency:

– 30 GHz is destriped with 0.25 s baselines; 44 and 70 GHz are destriped using 1 s baselines; and 100–857 GHz are destriped using pointing-period baselines (30–75 min).
– 30–70 GHz are destriped with a $1/f$-shape noise prior, while 100–857 GHz are destriped without a noise prior.
– 30, 44, and 70 GHz have separate destriping masks, while 100–857 GHz use the same 15% galaxy+point source mask.
– 30−70 GHz maps are destriped using baselines derived exclusively from the data going into the particular map, while 100−857 GHz maps are destriped using baselines derived from the full data set.

The baseline ($r = f_{NL} = 0$) fiducial channel/mission/full maps for temperature and total polarization are shown in Figs. 12 and 13, respectively.

4.2. Noise Monte Carlos

The FFP8 noise MCs are generated using `Madam/TOAST`, exploiting `TOAST`'s on-the-fly noise simulation capability to avoid the IO overhead of writing a simulated TOD to disk only to read it back in to map it (Fig. 14). In this implementation, `Madam` runs exactly as it would with real data, but whenever it submits a request to `TOAST` to provide it with the an interval of the noise TOD, that interval is simply simulated by `TOAST` in accordance with the noise power spectral densities provided in the runconfig, and returned to `Madam`.

For a simulation set of this size and complexity, requiring of the order of 10^{17} random numbers over 10^{12} disjoint and uncorrelated intervals, care must be take with the pseudo-random number generation to ensure that it is fast, reliable (and specifically uncorrelated), and reproducible, in particular enabling any process to generate any element of any subsequence on demand. To achieve this `TOAST` uses a Combined Multiple Recursive Generator (CMRG) (L'Ecuyer 1999) that provides more than sufficient period, excellent statistical robustness, and the ability to skip ahead to an arbitrary point in the pseudo-random sequence very quickly. Rather than seed each realization separately (with the associated risk of generating overlapping subsequences), we then set a single seed for the entire simulation set, divide the resulting sequence into 2^{64} disjoint subsequences, and assign a specific subset of subsequences to each pointing period of each detector for each realization, with the allocation of several rather than a single subsequence allowing us to support multiple independent timestream components, each of which depends on random number generation, such as the HFI detectors' uncorrelated and correlated noise components. Consequently when `Madam` requests the noise TOD for a particular interval of a particular

realization, the process can skip ahead to the appropriate location in the overall sequence; furthermore the random number generation itself can be threaded, with each thread then skipping ahead to the appropriate place within the subsequence.

`Madam` supports MC runs by enabling the pointing and flag data common to all realizations to be read once, before looping over the realizations that share them. However, to support maximum MC throughput some changes were also made to the standard `Madam` implementation:

– once the full data for any time span are loaded, `Madam/TOAST` can independently destripe and map the full and two half ring data sets;
– communication at high concurrency is more efficiently supported through use of the advanced MPI collective operation `alltoallv`;
– the map is distributed between the MPI processes based on pixels covered by the distributed TOD rather than using a simple round-robin scheme;
– HFI full-mission destriping is supported on-the-fly so that once the full mission data have been destriped, all subset maps can be binned as part of the same run;
– OpenMP/MPI hybridization allows for runs at higher concurrencies without suffering from MPI communication congestion.

4.3. CMB Monte Carlos

The FFP8 CMB MCs are generated using the `FEBeCoP` software package (Mitra et al. 2011), which produces beam-convolved maps directly in the pixel domain rather than sample-by-sample, as is done for the fiducial maps. The goal of this approach is to reduce the computational cost by the ratio of time-samples to map-pixels (i.e., the number of hits per pixel).

As illustrated in Fig. 15, `FEBeCoP` proceeds in 3 steps:

1. Given the satellite pointing and flags and the focal plane (accessed through the `TOAST` interface), for every channel `FEBeCoP` first re-orders all of the samples in the mission by pixel instead of time, localizing all of the observations of each pixel, and writes the resulting pixel-ordered detector dngles (PODA) to disk. Note that since the PODA also contains the detector, time-stamp, and weight of each observation this is a one-time operation for each frequency, and does not need to be re-run for different time intervals or detector subsets, or for changes in the beam model or its chosen cut-off radius.
2. For every time interval and detector subset to be mapped, and for every pixel in the map, `FEBeCoP` uses the PODA and the scanning beams to generate an effective-beam for that pixel which is essentially the weighted average of the discretized beam functions for every sample in the pixel included in the time interval and detector subset. The total effective-beam array is also written to disk. Given the PODA, this is a one-time operation for any beam definition.
3. Finally, `FEBeCoP` applies the effective-beam pixel-by-pixel to every CMB sky realization in the MC set to generate the corresponding beam-convolved CMB map realization.

The effective-beams provide a direct connection between the true and observed sky, explicitly incorporating the detailed pointing for every detector through a linear convolution. By providing the effective-beams at every pixel, `FEBeCoP` enables precise control of systematic effects, e.g., the point-spread functions
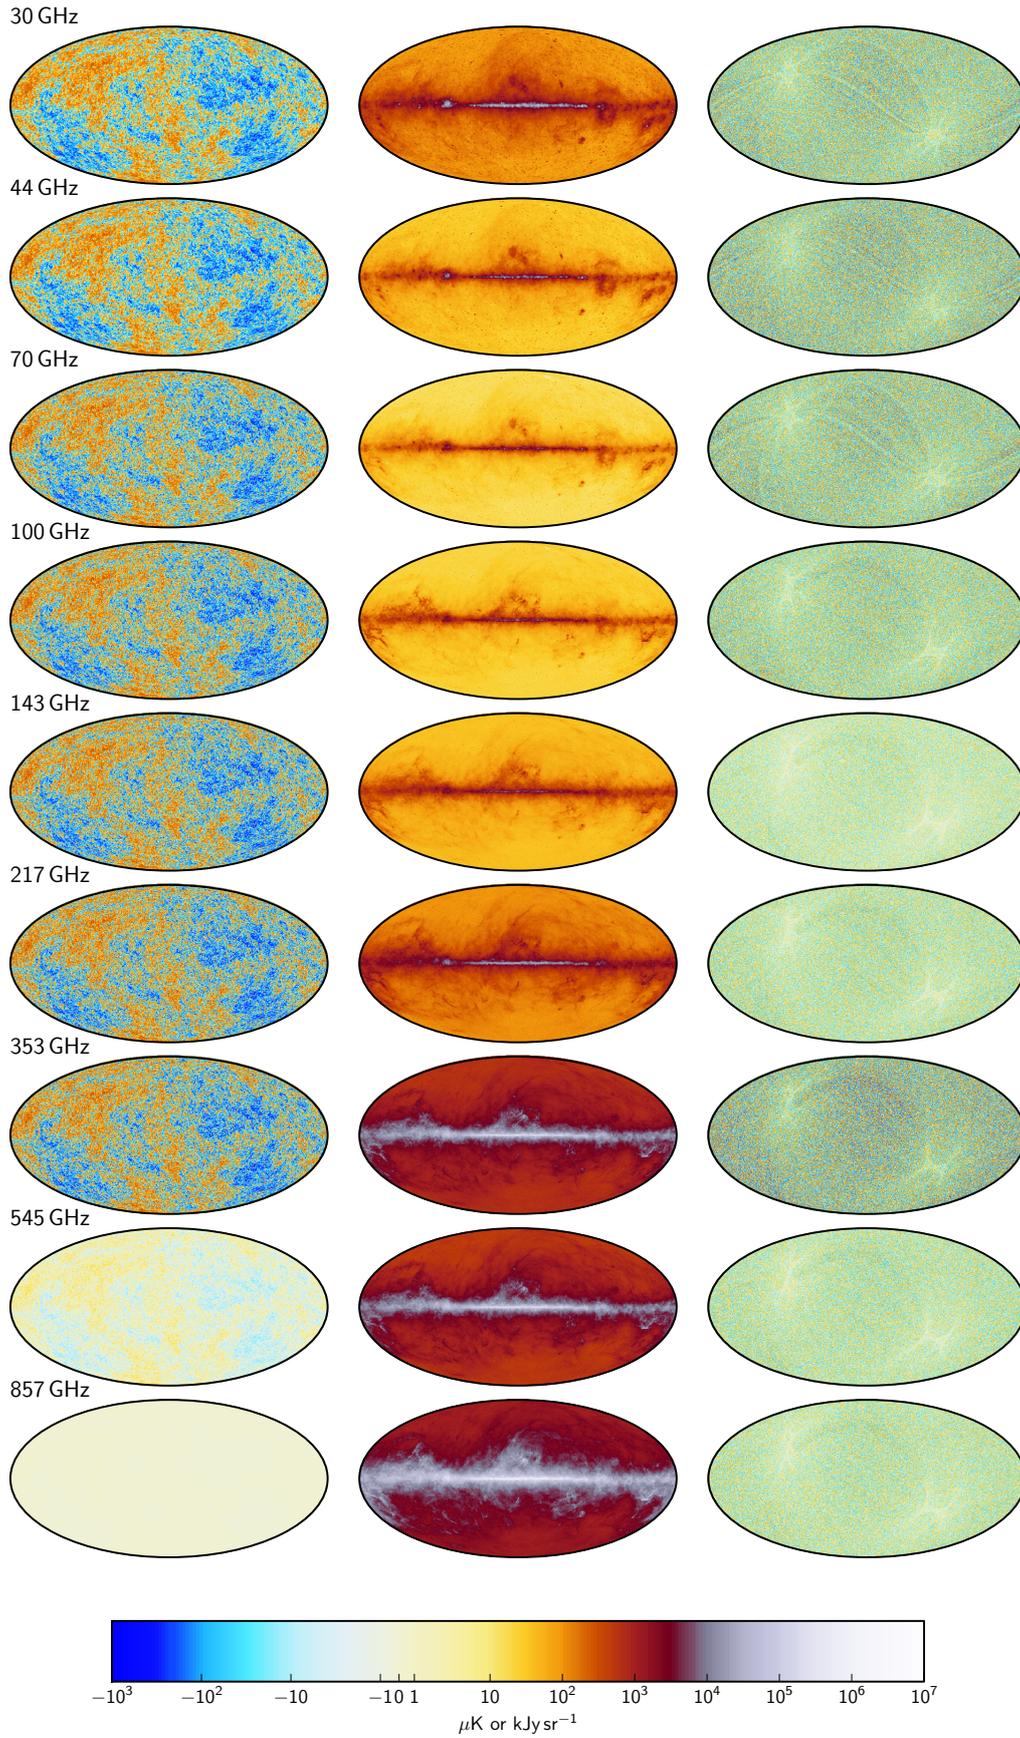
Fig. 12. *Left to right*: baseline FFP8 fiducial CMB, foreground and noise component temperature channel/mission/full maps at each frequency. Frequencies 30–353 GHz are plotted in μK, while 545 and 857 GHz are plotted in kJy sr^{-1}. See Fig. 16 for the combined temperature maps.
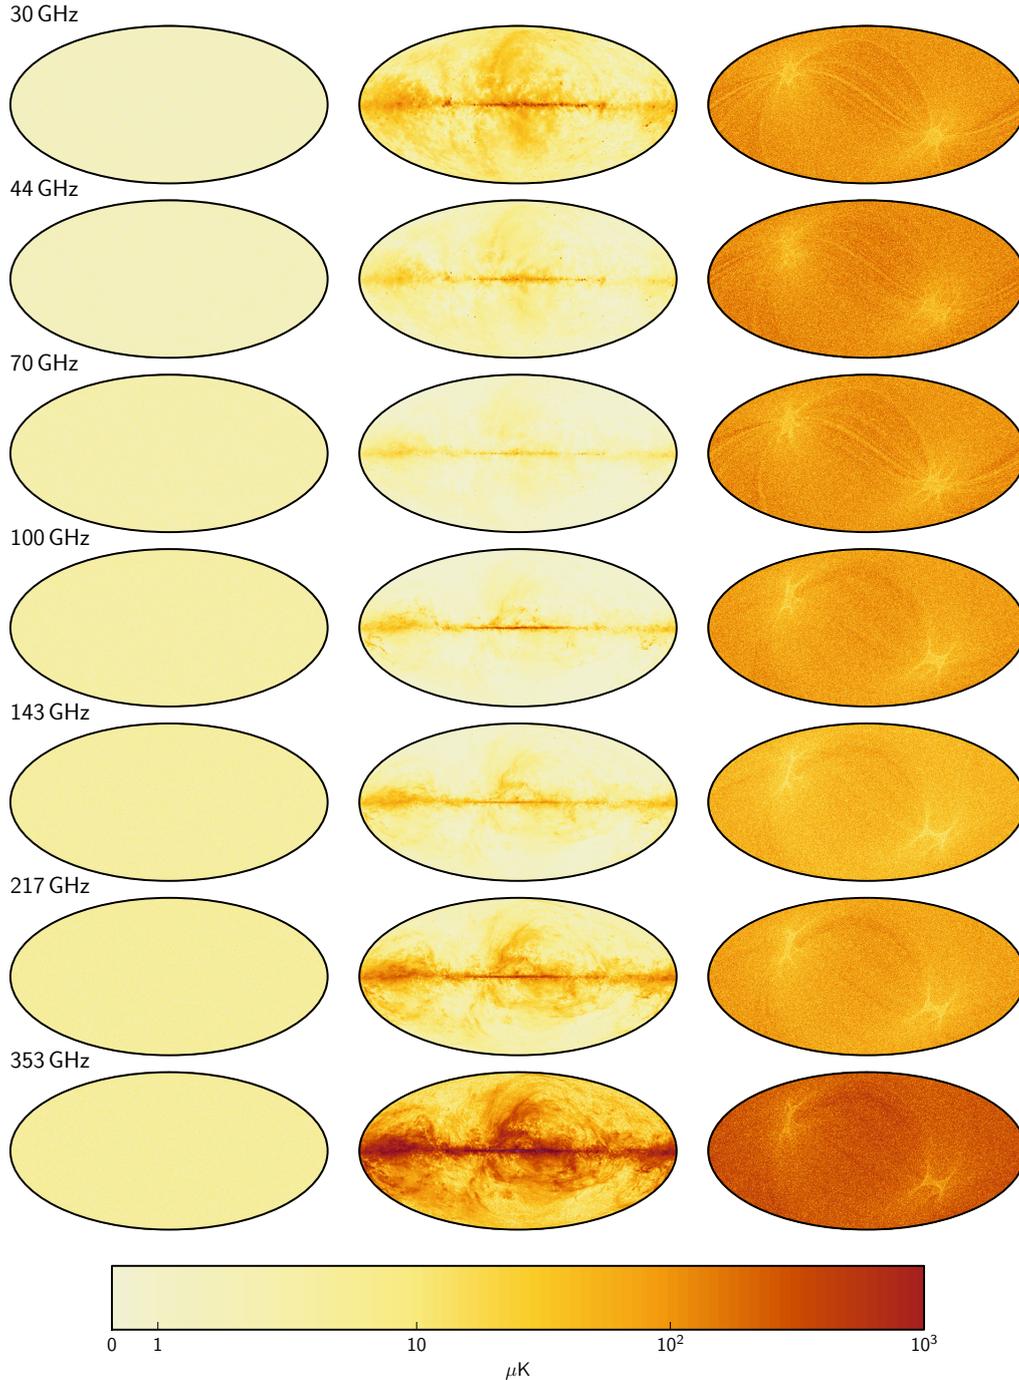
Fig. 13. *Left to right*: baseline FFP8 fiducial CMB, foreground and noise component total polarization channel/mission/full maps at each frequency. See Fig. 17 for the combined polarization maps.

can be fitted at each pixel on the sky and used to determine point source fluxes (Planck Collaboration XXVI 2016; Planck Collaboration XXVII 2016).

4.4. Validation

Our goal for the FFP8 simulation set is that it be not only internally self-consistent, but also a good representation of the real data. In addition to the validation steps carried out on all of the inputs individually and noted in their respective sections above, we must also validate the final outputs. A first crude level of validation is provided simply by visual inspection of the FFP8 and

real *Planck* maps – Figs. 16 (temperature) and 17 (total polarization) – where the only immediately apparent difference is the CMB realization.

While this is a necessary test, it is hardly sufficient, and the next step is to compare the angular power spectra of the simulated and real channel/mission/full maps. As illustrated in Fig. 18 and Table 5, LFI channels show excellent agreement across all angular scales, while HFI channels show a significant power deficit at almost all angular scales. Since this missing HFI power is not picked up in the noise estimation, it must be sky-synchronous (frequency bins corresponding to sky-synchronous signals being discarded when fitting the noise PSDs due to
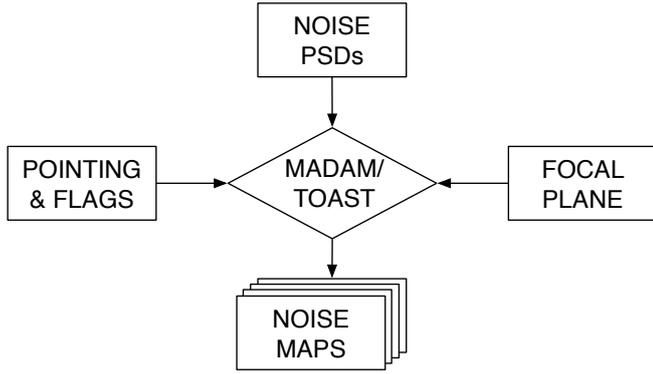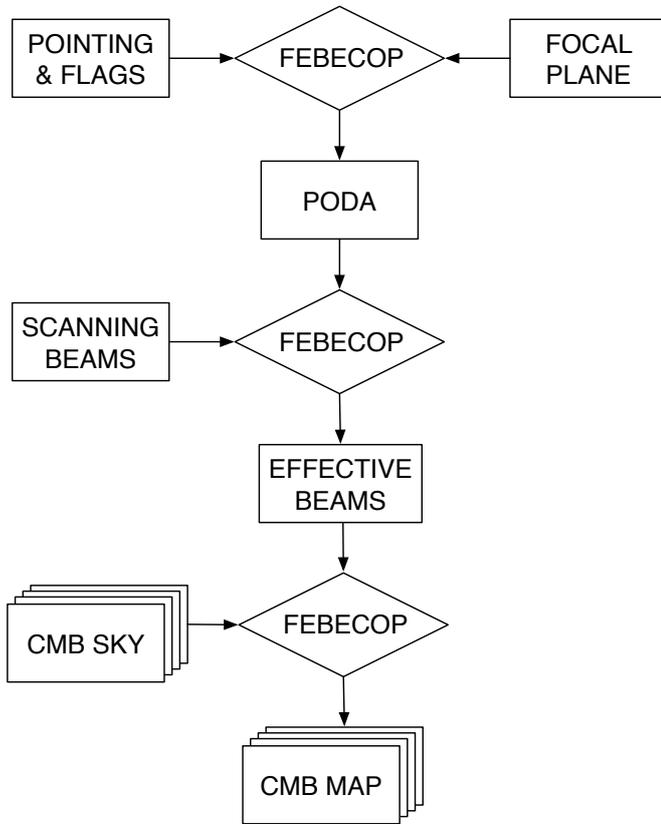
Fig. 14. Noise Monte Carlo pipeline schematic.

Fig. 15. CMB Monte Carlo pipeline schematic.

Table 5. *EE* and *BB* mission full map spectrum ratios.

Frequency [GHz]	*EE*	*BB*
30	0.998	1.000
44	0.999	1.000
70	0.998	1.000
100	1.019	1.020
143	1.033	1.029
217	1.017	1.007
353	1.042	1.039

This is the full mission median C_ℓ ratio (PR2-2015 / FFP8) in the noise-dominated range $\ell \in [2000, 3000]$ (LFI) and $\ell \in [3000, 4000]$ (HFI); as expected, FFP8.1 exhibits the same values.

found in the relevant papers (e.g., Planck Collaboration IX 2016; Planck Collaboration X 2016; Planck Collaboration XI 2016; Planck Collaboration XVI 2016).

5. High performance computing considerations

5.1. Simulations at NERSC

NERSC is the US Department of Energy's general purpose supercomputing centre, with a new top-10 HPC system[14] being installed about every 3 years. For the last 15 years it has provided annual allocations of cycles and storage to the *Planck* collaboration, with a unique agreement between NASA and DOE guaranteeing *Planck* access to NERSC through the mission lifetime (Borrill et al. 2014). The US Planck Collaboration has also purchased exceptional levels of service at NERSC, including 240 TB of globally accesssible disk space and one cabinet (640 cores) of the 9984-core IBM iDataplex *Carver*. In 2014, *Planck* was awarded about 75 000 000 NERSC CPU-hours, which were shared by the 170 *Planck* data analysts who have requested NERSC accounts.

For FFP8, the fiducial realization, all HFI noise MCs, the first 1000 LFI channel noise MCs, and all CMB MCs were performed on NERSC's 133 824-core Cray XC30 *Edison*. Subsequently, all of the FFP8.1 CMB MCs were generated there too.

5.1.1. The fiducial realization

The FFP8 fiducial foreground and CMB components were generated via explicit timelines and the noise via an implicit timeline, as detailed in the noise MC section below. For the explicit cases each frequency used a separate job script, allowing multiple frequencies to be executed simultaneously without concerns about load-balancing or synchronization. Within each frequency-specific job each component was run in turn, allowing the same target TOD files to be re-used (overwritten) to save disk space. Within each component, each detector was run in turn to prevent overloading the file system (since the exchange format TODs are stored by frequency). Once a full set of TODs had been generated for a given component at a given frequency, the full set of 1134 maps was produced before moving on to the next component. The concurrencies and runtimes for each frequency are listed in Table 6. With five components generated via explicit timelines and one via on-the-fly simulation within the map-making, the total cost was approximately 250 000 CPU-hours. This same set-up also allowed us to generate each individual

their contamination by signal residuals). This is now understood to be a systematic effect introduced in the HFI pre-processing pipeline, and we are working both to incorporate it as a systematic component in existing simulations and to ameliorate if for future data releases.

Finally, the various analyses of the FFP8 maps in conjuction with the flight data provide powerful incidental validation. To date the only issues observed here are the known mismatch between the FFP8 and PR2-2015 cosmologies, and the missing systematic component in the HFI maps. As noted above, the former is readily addressed by rescaling or using FFP8.1; however, the characterization and reproduction of the latter is an ongoing effort. Specific details of the consequences of this as-yet-unresolved issue, such as its impact on null-test failures and *p*-value stability in studies of non-Gaussianity, can be
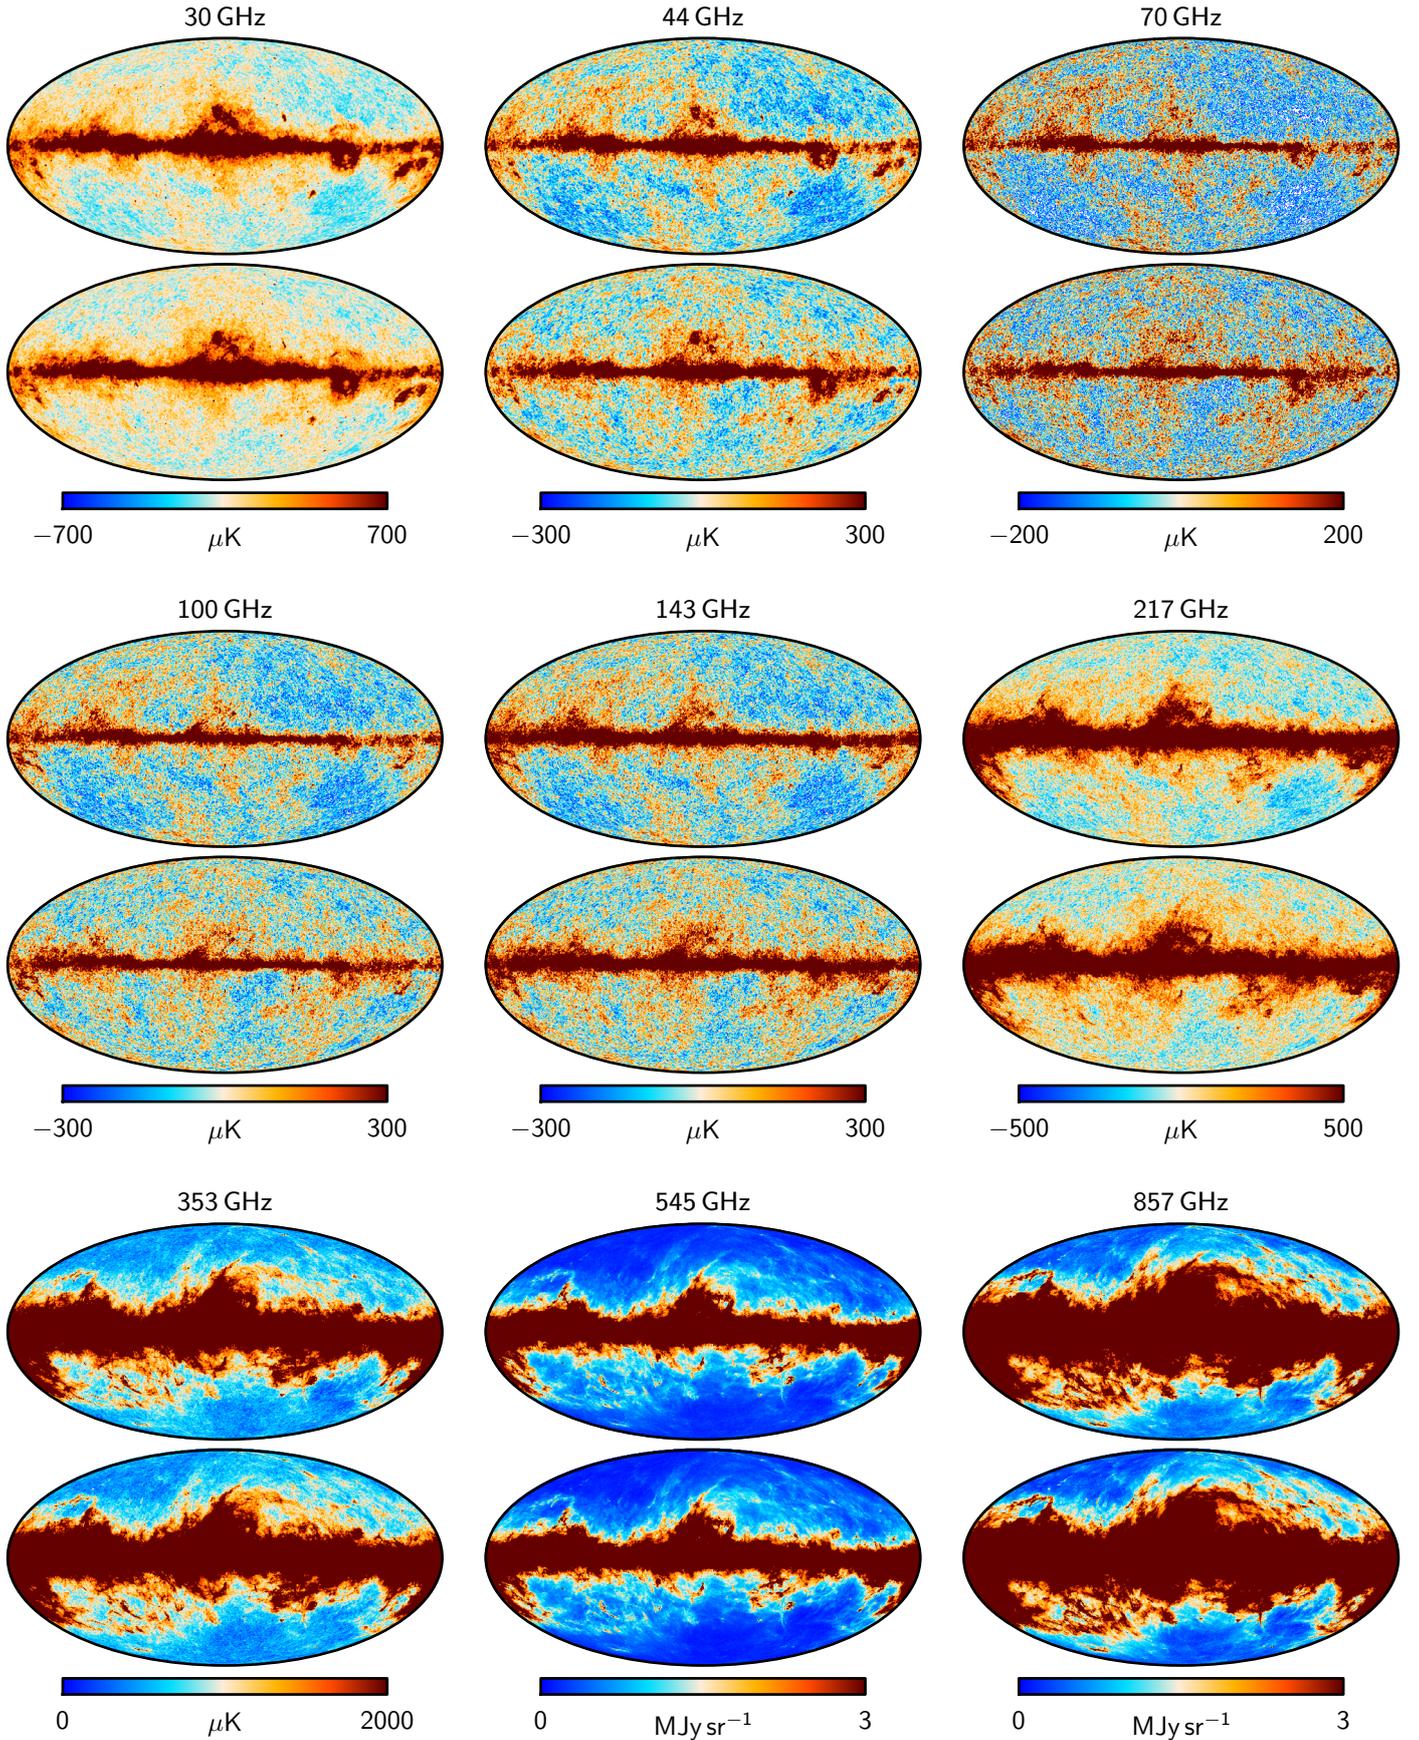
---

[14] http://top500.org

Fig. 16. Pairs of channel/mission/full temperature maps, comparing FFP8 simulations (*upper*) and 2015 *Planck* data (*lower*). All maps are downgraded to $N_{\mathrm{side}} = 256$. Visible differences away from the galactic plane are between the actual and simulated CMB sky.
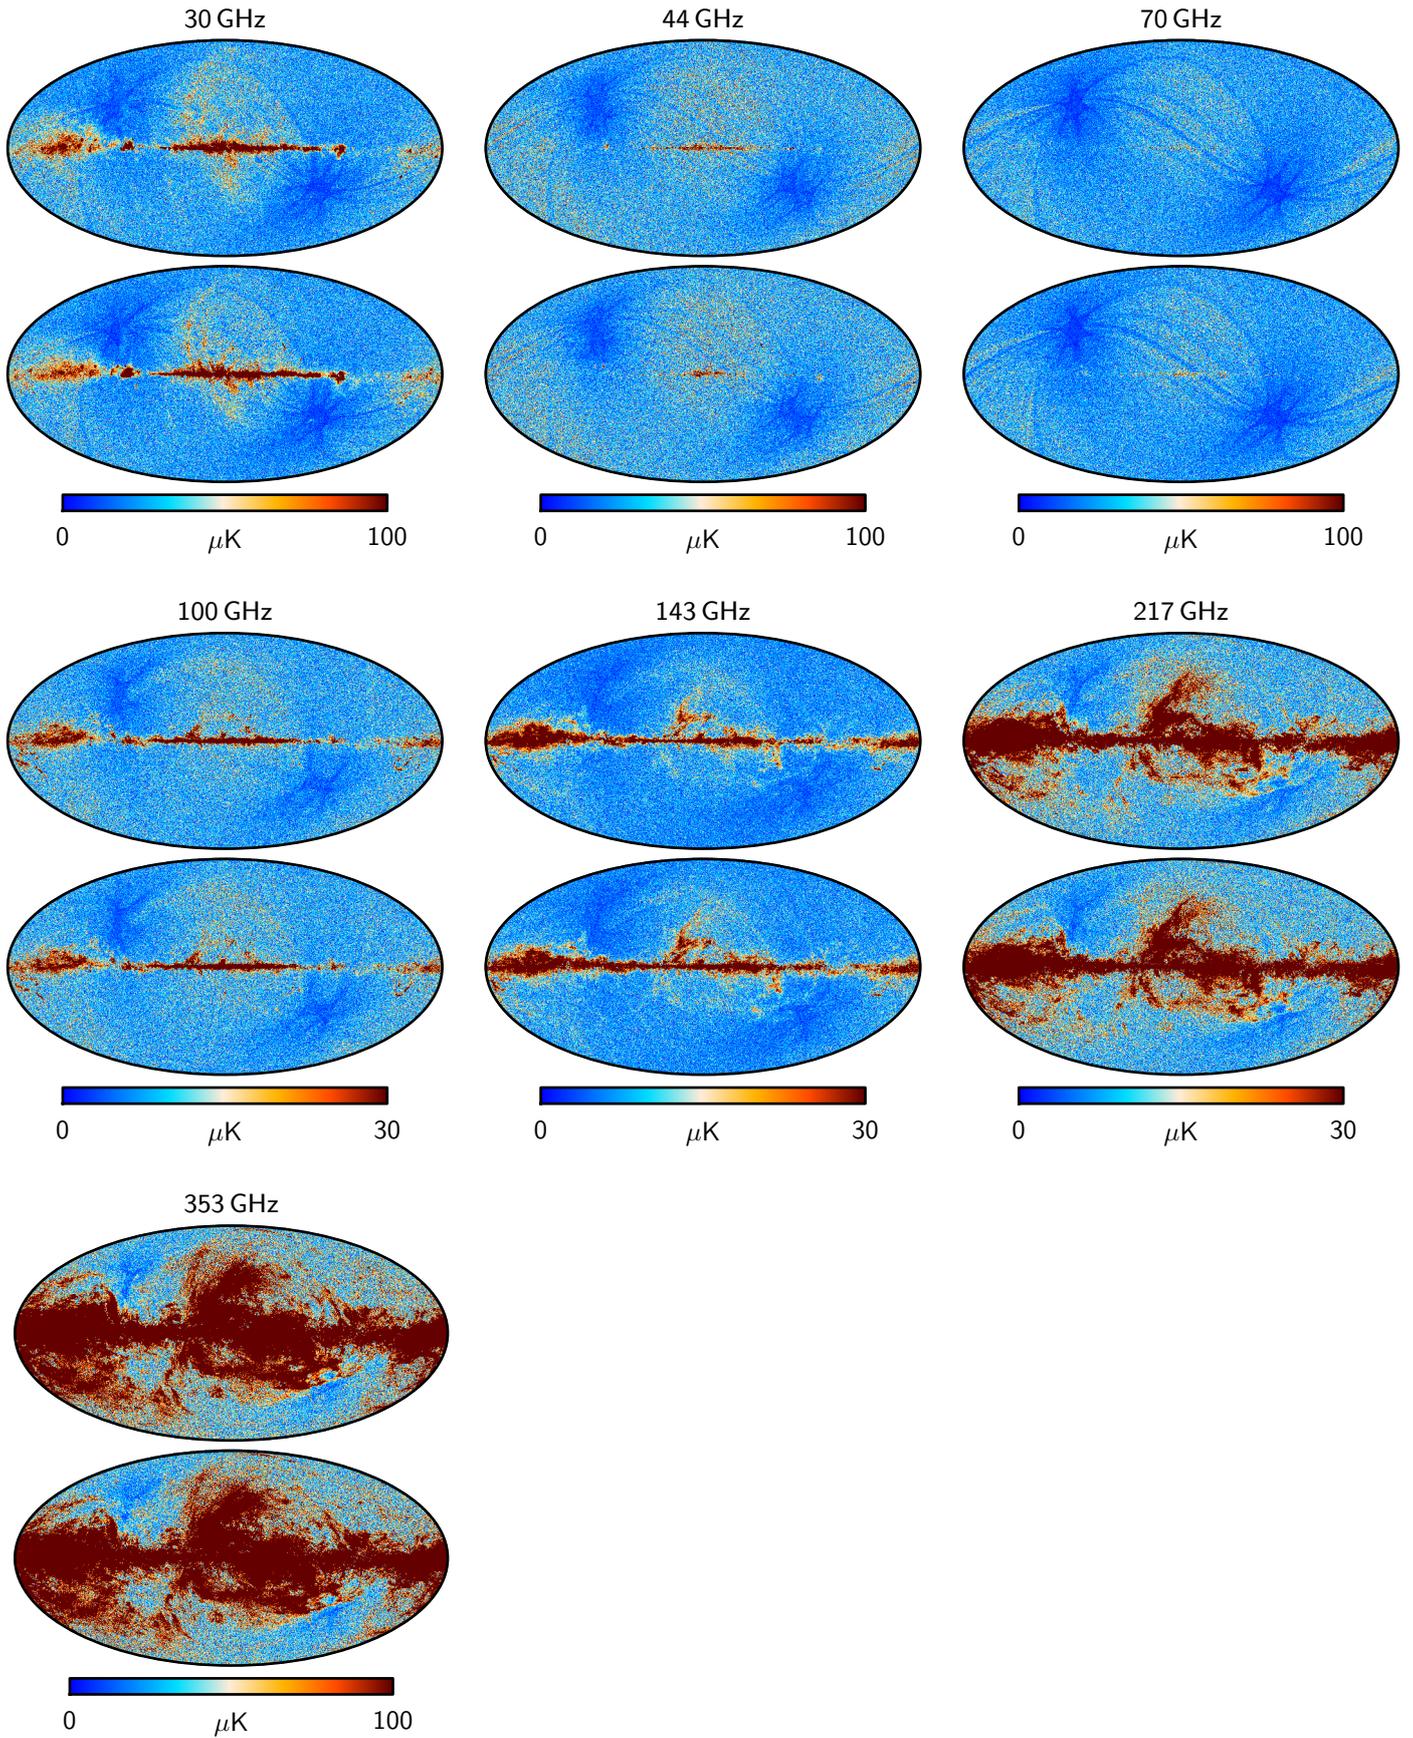
Fig. 17. Pairs of channel/mission/full total polarization maps, comparing FFP8 simulations (*upper*) and 2015 *Planck* (*lower*) data. All maps are downgraded to $N_{side} = 256$.
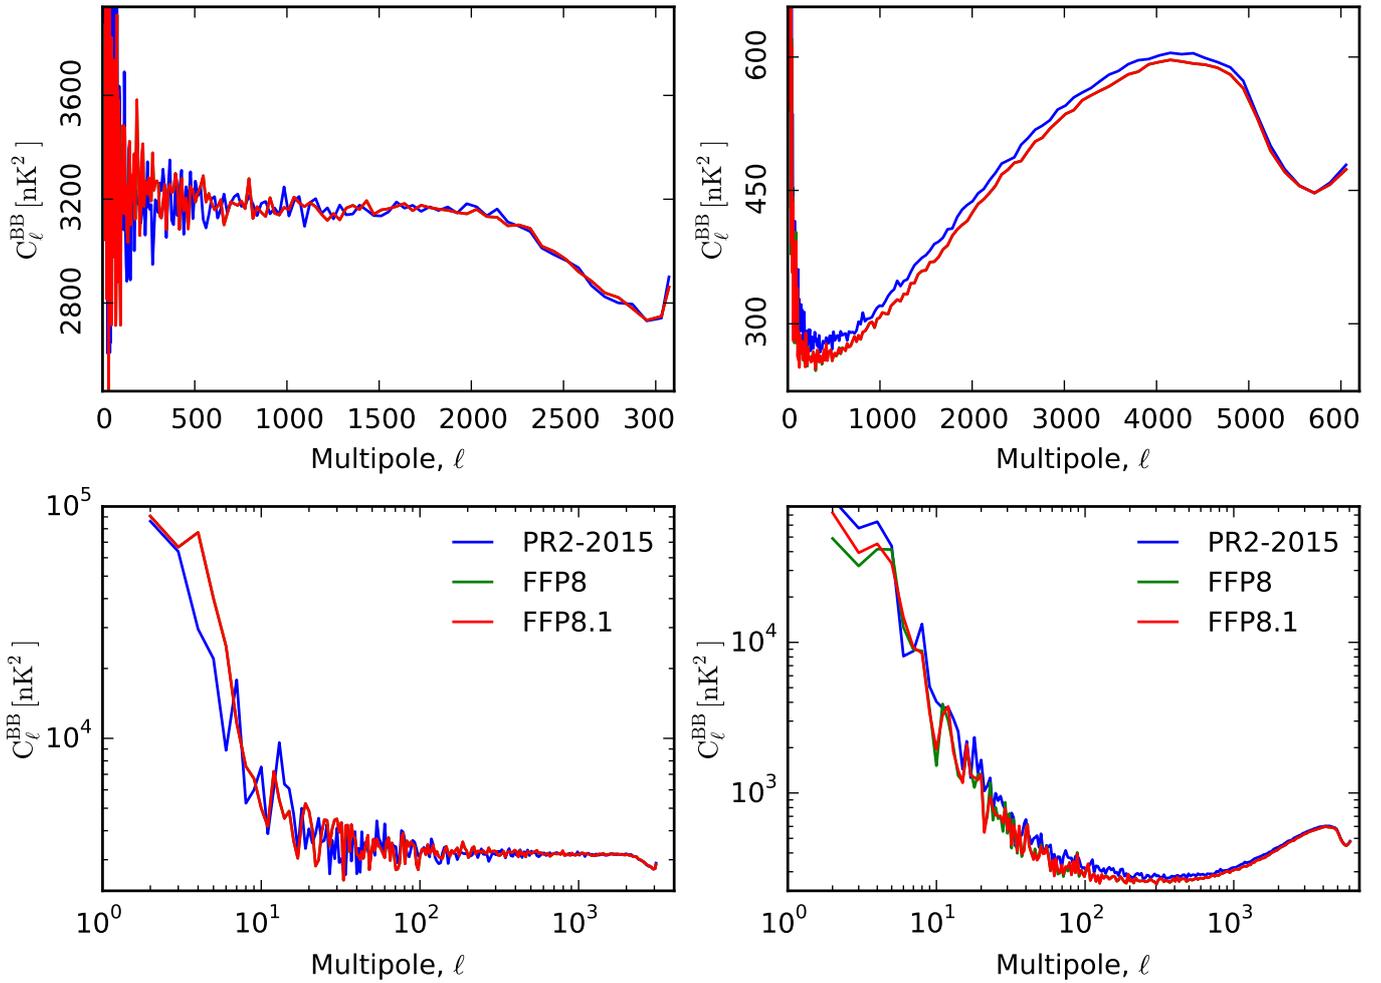
Fig. 18. Comparison of the *BB* spectra of the channel/mission/full simulated and flight data maps using pseudo-spectra computed on 75% of the sky with the Galaxy and point sources masked, showing excellent agreement at 70 GHz but a few percent discrepancy at almost all angular scales at 100 GHz. *Left*: 70 GHz full map *BB*-spectra. *Right*: 100 GHz full map *BB*-spectra. *Top*: linear horizontal axis shows small scale behaviour. *Bottom*: logarithmic horizontal axis shows large scale behaviour.

Table 6. Computational cost of generating one component of the FFP8 fiducial realization for each frequency.

Frequency [GHz]	Concurrency (Cores)	TODs [min]	Maps [min]	Cost [CPU-hrs]
30	3600	12	20	1920
44	3600	36	15	3060
70	3600	108	48	9360
100	4800	64	3	5360
143	4800	88	5	7440
217	4800	96	5	8080
353	4800	96	5	8080
545	4800	18	2	1600
857	4800	24	2	2080

Notes. TOD generation includes all of the detectors, and mapmaking all of the maps, at that frequency. Note the much higher cost of LFI mapmaking, due to its shorter baselines and run-baseline destriping.

foreground component to cross-check the component separation pipelines.

5.1.2. The noise Monte Carlos

For the noise MCs, we noted that mapmaking runs most efficiently at the lowest concurrency with sufficient memory to hold all of the data at a given frequency. Additional parallelism was then achieved by running multiple instances of a single frequency simultaneously. In practice, for each instrument and frequency we first ran a job containing 10 instances, each generating 10^2 realizations of all map-types, to produce the first 10^3 realizations, and then, for HFI only (for LFI, see Sect. 5.2), ran a second job running nine instances each generating 10^3 realizations of only the channel/mission/full maps to fill out the full 10^4 realizations. For LFI, where every data subset is individually destriped, the instances were further subdivided to keep the data per core constant, so that the same number of cores would process one channel/mission, two channel/half-mission, four channel/year, or eight channel/survey maps, or (for the 70 GHz data) three detset/mission, six detset/half-mission, twelve detset/year or twenty-four detset/survey maps. Table 7 shows the cost per realization of all the maps at one frequency (for realizations 0−999) and of the HFI channel/mission/full maps alone (for realizations 1000−9999). The overall cost of these runs was approximately

Table 7. Computational cost of generating one realization of the FFP8 noise MC for each frequency, both for the full set of maps and for the HFI channel/mission/full (CMF) maps alone.

Frequency [GHz]	Concurrency (Cores)	All maps		CMF maps	
		[min]	[CPU-hrs]	[min]	[CPU-hrs]
30	2304	10	384	. . .	. . .
44	3456	8	460	. . .	. . .
70	4608	13	949	. . .	. . .
100	2640	3	176	2	88
143	3600	3	180	2	120
217	3960	3	198	2	132
353	3960	3	198	2	132
545	600	4	40	2	20
857	600	6	60	3	30

Table 8. Numbers of samples ($\mathcal{N}_t$), map-pixels ($\mathcal{N}_p$), beam-pixels ($\mathcal{N}_b$) and derived computational costs of the FEBeCoP effective-beam calculation ($\mathcal{N}_t \mathcal{N}_b$) and MC convolution ($\mathcal{N}_{mc} \mathcal{N}_p \mathcal{N}_b$) for each channel in FFP8.

Frequency	$\mathcal{N}_t$	$\mathcal{N}_p$	$\mathcal{N}_b$	$\mathcal{N}_t \mathcal{N}_b$	$\mathcal{N}_{mc} \mathcal{N}_p \mathcal{N}_b$
30 . . .	1.5×10^{10}	3.8×10^6	3.5×10^3	5.2×10^{13}	1.3×10^{15}
44 . . .	3.2×10^{10}	3.8×10^6	1.7×10^3	5.4×10^{13}	6.4×10^{14}
70 . . .	1.1×10^{11}	3.8×10^6	7.3×10^2	7.9×10^{13}	2.8×10^{14}
100 . . .	7.5×10^{10}	1.5×10^7	1.1×10^4	8.3×10^{14}	1.7×10^{16}
143 . . .	1.0×10^{11}	1.5×10^7	1.1×10^4	1.1×10^{15}	1.7×10^{16}
217 . . .	1.1×10^{11}	1.5×10^7	1.1×10^4	1.2×10^{15}	1.7×10^{16}
353 . . .	1.1×10^{11}	1.5×10^7	1.1×10^4	1.2×10^{15}	1.7×10^{16}
545 . . .	3.2×10^{10}	5.0×10^6	1.1×10^4	3.6×10^{14}	5.5×10^{15}
857 . . .	4.3×10^{10}	5.0×10^6	1.1×10^4	4.8×10^{14}	5.5×10^{15}

Notes. The number of effective-beam pixels at each frequency depends on the HEALPix resolution of the map (N_{side} = 2048 and 1024 for HFI and LFI, respectively) and the cut-off radius of the scanning beam (113.6′, 79′, 52′ and 100′ for 30 GHz, 44 GHz, 70 GHz and HFI, respectively).

3 000 000 CPU-hours for the 10^3 realizations of all maps and 5M CPU-hours for the 9×10^3 additional HFI channel/mission/full maps.

5.1.3. The CMB Monte Carlos

For the CMB MCs, the computational cost depends on the numbers of time-samples $\mathcal{N}_t$, map-pixels $\mathcal{N}_p$, beam-pixels $\mathcal{N}_b$, and MC realizations $\mathcal{N}_{mc}$. Generating the PODA is essentially a re-ordering of the time-domain data, scaling as $\mathcal{N}_t$; generating the effective-beam is a calculation of beam-pixel weights for every observation, scaling as $\mathcal{N}_t \mathcal{N}_b$; and generating the beam-convolved maps is an effective-beam-matrix/map-vector multiplication for each realization, scaling as $\mathcal{N}_{mc} \mathcal{N}_p \mathcal{N}_b$. Table 8 shows the values of these parameters for each *Planck* channel, indicating that for FFP8 the effective-beam convolution over the CMB skies has, for the first time, become the dominant cost. By comparison, for FFP6, although the numbers of samples were a factor of 2 (HFI) to 4 (LFI) smaller, the smaller MC set (with only 10^3 realizations) and smaller numbers of pixels (temperature only) made the effective-beam calculation dominate. Running the 46 distinct 10^4-realization CMB MC map sets included in FFP8 then required some 8M CPU-hours in total.

Originally FEBeCoP was designed to optimize MC simulations with well-localized beams (and hence small $\mathcal{N}_b$), but it is clear that with FFP8 we have moved out of that regime and it is time to re-evaluate the comparison of the computational cost and accuracy of the effective-beam and explicit time-domain approaches. Since these scale as $\mathcal{N}_p \mathcal{N}_b$ and $\mathcal{N}_t$, respectively, the cross-over point is when the number of pixels in the effective-beam is comparable to the number of hits per map-pixel. Possible ways to reduce $\mathcal{N}_b$ for FEBeCoP would be to use multi-component beams with each component pixelized at a different resolution, or to restrict the main beam to the pixel domain and process the intermediate and far sidelobes in the spherical harmonic domain. Looking forward, this will be an active area of research and development. We note, though, that next-generation *B*-mode experiments are specifically designed to have very many more hits per pixel than *Planck* in order to achieve the necessary signal-to-noise ratio to meet their science goals, shifting the balance back towards pixel-domain calculations.

5.2. Simulations at CSC

A significant fraction of the LFI noise MCs were performed at the CSC-IT Center for Science in Finland, using their Cray XC30 supercomputer *Sisu* with 40 512 computing cores. These runs were performed using the same software stack as was used at NERSC, with cross-checks confirming that the two systems produce results identical to numerical precision.

The CSC runs included the channel/mission/full and half-ring maps for realizations 1000–9999, totalling 81 000 maps and using a total of 4 million CPU-hours. These were run in about a dozen different jobs, gradually increasing the number of instances of the code running simultaneously and the number of cores used. The largest jobs for each frequency were as follows:

– for 30 GHz, 10 simultaneous instances generating 5000 realizations in all, running on 23 040 cores for 17.8 h for a total of 410 000 CPU-hours;

– for 44 GHz, 10 simultaneous instances generating 5000 realizations in all, running on 11 520 cores for 21.8 h for a total of 251 000 CPU-hours;

– for 70 GHz, 10 simultaneous instances generating 2757 realizations in all, running on 36 000 cores for 24 h[15] for a total of 864 000 core-hours.

In addition, we produced a smaller number of realizations for the 70 GHz detector subsets, including 1000 realizations of the mission and 300 realizations each of the two half-missions, four years, and eight surveys. With three detector subsets and three maps (full and half-ring) for every realization, this resulted in 46 800 maps and took a total of 200 000 CPU-hours.

In total, 127 800 maps were produced at CSC, using 4 200 000 CPU-hours and requiring a total of 17.7 TB of disk space; these were then transferred to NERSC for wider distribution.

6. Results

The core FFP8 simulation set consists of:

– 90 720 fiducial maps, comprising 1134 maps of each of six components (three CMB, two foreground, one noise), then

---

[15] The maximum runtime allowed at CSC for jobs of this size.

combined to produce two sets of total maps (with and without bandpass mismatch), all for five different values of the (r, f_{NL}) parameter pair.

- 671 400 noise MC maps, comprising 10^4 realizations of all channel/mission/full maps, at least 10^3 realizations of all other full maps, and at least 10^2 realizations of all half-ring maps.
- 460 000 CMB MC maps, comprising 10^4 channel/mission/full and channel/half-mission/full maps and selected detset/mission/full maps.

while the FFP8.1 simulation set currently consists of:

- 18 144 fiducial maps, comprising 1134 maps of each of three CMB components, then combined with the FFP8 foreground and noise maps to produce two sets of total maps (with and without bandpass mismatch), restricted the baseline ($r = f_{NL} = 0$) case.
- 27 000 CMB MC maps, comprising 1000 channel/mission/full and channel/half-mission/full maps.

6.1. Uses of FFP8

As well as supporting the validation and verification of all of the *Planck* analysis codes, FFP8 is widely used in the analysis of *Planck* data.

Fiducial realizations are used to:

- test how various choices of the smoothing kernel affect signal aliasing in low resolution maps (LFI) (Planck Collaboration VI 2016);
- demonstrate inaccuracies in the measurements of the instrumental bandpasses (Planck Collaboration X 2016);
- measure and subtract the LFI sidelobe pick-up (Planck Collaboration II 2016);
- estimate calibration and bandpass measurement uncertainties in component separation (Planck Collaboration X 2016);
- de-bias and measure uncertainty in lensing reconstruction (Planck Collaboration XV 2016);
- assess SZ catalog reliability (Planck Collaboration XXVII 2016);
- estimate noise for mapping the thermal Sunyaev-Zeldovich (tSZ) effect (Planck Collaboration XXII 2016).

Noise MCs are used to:

- measure the accuracy of the LFI pixel-pixel noise covariance matrix construction (Planck Collaboration VI 2016); and
- inform low-ℓ data selection by measuring the distribution of low-multipole amplitudes due to noise as well as confirming the anomalous nature of Surveys 2 and 4 (Planck Collaboration II 2016);

CMB MCs are used to:

- produce beam transfer functions for each map type listed in Table 1 by averaging over the corresponding MC realizations (Planck Collaboration IV 2016; Planck Collaboration VII 2016).

In addition, the noise and CMB MC maps are combined using the same component separation pipelines as the real and fiducial data. The resulting MC sets are then used to support the various CMB isotropy and statistics analyses

(Planck Collaboration XVI 2016) and non-Gaussianity estimators (Planck Collaboration XVII 2016).

Various ancilliary data sets were also produced to address the specific requirements of particular analyses. The most notable of these were:

- nine channel/mission/full maps for each of the ten individual foreground components, to validate the component separation analyses;
- 10^3 channel/mission/full of the scalar CMB MCs with no Doppler boosting applied, to validate the `BipoSH` reconstruction of Doppler boosting (Planck Collaboration XVI 2016).

The capabilities developed for FFP8 also allowed for many other tests of data and analysis systematics, including running the mapmaking:

- with both mission- and run-baseline destriping to measure correlations induced by mission-baseline destriping in maps spanning different time intervals or including disjoint detector subsets, leading to the acceptance of this approach by HFI (where these are generally small) and its rejection by LFI (where they are significant);
- with both temperature-only and temperature+polarization destriping to demonstrate the loss of polarization power induced by the former, leading to the adoption of the latter by both HFI and LFI;
- using pointing that includes the expected uncertainty in its reconstruction to test the impact of pointing reconstruction error for LFI maps;
- using polarization angles that include the expected uncertainty in their determination to test the impact of polarization angle error for LFI maps.

6.2. Known issues with FFP8

Generating a simulation set of the size and complexity of FFP8 requires the coordination of a large number of people and processes. In addition to the cosmology mismatch and HFI systematic residual issues noted above, there are a few errors that are known to have been inadvertently introduced into FFP8 during its execution.

- Due to an approximation in the way the PSM applies the detector bandpasses to the input skies the effective bandpasses differ from the inputs by up to 2% in the submillimetre channels. This effect has been incorporated into the fiducial maps, but when the submillimetre CMB MC maps are used with the fiducial realization they must be rescaled by factors of 0.9976 at 545 GHz and 0.9828 at 857 GHz.
- Due to a bug in the code generating the CMB realizations, the aberration contribution to the Doppler boosting was inadvertently omitted. New, corrected, realizations are being generated and will be added to the public data release when they become available.
- During the export of the TOD from the HFI DPC the sample time stamps were inadvertently truncated, resulting in pointing errors of up to $0\rlap{.}'3$.
- The channel-average bandpasses generated by the PSM used preliminary detector weights, resulting in a tiny discrepancy between the bandpass-mismatch and no-bandpass-mismatch cases.
- The FFP8 mapmaking used a slightly more rigorous pixel-rejection criterion than the HFI DPC, resulting in some FFP8

maps having somewhat more missing pixels than the real data. Since these are the noisiest pixels, their naive inclusion in any map-based analysis can have a disproportionate effect at high ℓ.

6.3. Community access to FFP8

As was the case with FFP6 for the first *Planck* data release, FFP8/FFP8.1 is being made available to the scientific community. A simulation set of this size requires significant computational resources to store and analyse, so once again we are making it available at NERSC, where there is also an allocation of compute cycles available to support its exploitation. Full details on this, including how to obtain a NERSC account, can be found online[16].

7. Conclusions

Simulations play a critical role in the analysis of CMB data sets, most particularly in the validation and verification of analysis pipelines and the debiasing and uncertainty quantification of analysis results. The former requires a simulation that is as true to the mission as possible, while the latter requires enormous Monte Carlo sets of mission realizations. FFP8 provides the 2015 *Planck* data release with both. The fiducial realization takes the most realistic model of the microwave sky available to us and simulates the full mission observation of that sky for every *Planck* detector, incorporating our best estimates of that detector's beam, bandpass, and time-varying noise properties. The noise and CMB MC sets include 10^4 realizations of the most important maps (namely those that incorporate all of the data at a given frequency), together with 10^2–10^4 realizations of all of the other maps that include subsets of the observations and detectors at each frequency. The FFP8.1 CMB maps provide a smaller sample of fiducial and MC CMB skies consistent with the PR2-2015 cosmology. With around 1.25 million maps in total, and requiring some 25 million CPU-hours at multiple HPC centres, this is by far the largest CMB simulation set ever produced.

References

Akaike, H. 1974, IEEE Transactions on Automatic Control, 19, 716
Arnaud, M., Pointecouteau, E., & Pratt, G. W. 2005, A&A, 441, 893
Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A&A, 517, A92
Babich, D., Creminelli, P., & Zaldarriaga, M. 2004, JCAP, 8, 9
Bennett, C. L., Smoot, G. F., Hinshaw, G., et al. 1992, ApJ, 396, L7
Bennett, C. L., Hill, R. S., Hinshaw, G., et al. 2003, ApJS, 148, 97
Bersanelli, M., Mandolesi, N., Butler, R. C., et al. 2010, A&A, 520, A4
Bertschinger, E. 1987, ApJ, 323, L103
Béthermin, M., Wang, L., Doré, O., et al. 2013, A&A, 557, A66
Blas, D., Lesgourgues, J., & Tram, T. 2011, J. Cosmol. Astropart. Phys., 7, 34
Borrill, J., Keskitalo, R., & Kisner, T. 2014, CISE, 17, 22

Cai, Z.-Y., Lapi, A., Xia, J.-Q., et al. 2013, ApJ, 768, 21
Challinor, A., & van Leeuwen, F. 2002, Phys. Rev. D, 65, 103001
Chen, C.-C., Cowie, L. L., Barger, A. J., et al. 2013, ApJ, 762, 81
Davies, R. D., Dickinson, C., Banday, A. J., et al. 2006, MNRAS, 370, 1125
Delabrouille, J., Melin, J.-B., & Bartlett, J. G. 2002, in AMiBA 2001: High-Z Clusters, Missing Baryons, and CMB Polarization, eds. L.-W. Chen, C.-P. Ma, K.-W. Ng, & U.-L. Pen, ASP Conf. Ser., 257, 81
Delabrouille, J., Betoule, M., Melin, J.-B., et al. 2013, A&A, 553, A96
de Oliveira-Costa, A., Tegmark, M., Gutiérrez, C. M., et al. 1999, ApJ, 527, L9
de Zotti, G., Massardi, M., Negrello, M., & Wall, J. 2010, A&ARv, 18, 1
Dickinson, C., Davies, R. D., & Davis, R. J. 2003, MNRAS, 341, 369
Draine, B. T., & Lazarian, A. 1998, ApJ, 508, 157
Elsner, F., & Wandelt, B. D. 2009, ApJS, 184, 264
Enqvist, K., & Sloth, M. S. 2002, Nucl. Phys. B, 626, 395
Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, ApJ, 524, 867
González-Nuevo, J., Massardi, M., Argüeso, F., et al. 2008, MNRAS, 384, 711
Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, ApJS, 180, 225
Itoh, N., Kohyama, Y., & Nozawa, S. 1998, ApJ, 502, 7
Kamionkowski, M., & Knox, L. 2003, Phys. Rev. D, 67, 063001
Keihänen, E., Kurki-Suonio, H., & Poutanen, T. 2005, MNRAS, 360, 390
Keihänen, E., Keskitalo, R., Kurki-Suonio, H., Poutanen, T., & Sirviö, A. 2010, A&A, 510, A57
Koester, B. P., McKay, T. A., Annis, J., et al. 2007, ApJ, 660, 239
Kogut, A., Banday, A. J., Bennett, C. L., et al. 1996, ApJ, 464, L5
Kogut, A., Fixsen, D. J., Levin, S. M., et al. 2011, ApJ, 734, 4
Lamarre, J., Puget, J., Ade, P. A. R., et al. 2010, A&A, 520, A9
L'Ecuyer, P. 1999, Operations Res., 47, 159
Leitch, E. M., Readhead, A. C. S., Pearson, T. J., & Myers, S. T. 1997, ApJ, 486, L23
Lewis, A. 2005, Phys. Rev., D71, 083008
Lewis, A. 2013, J. Cosmol. Astropart. Phys., 1308, 053
Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473
Liguori, M., Matarrese, S., & Moscardini, L. 2003, ApJ, 597, 57
Lyth, D. H., Ungarelli, C., & Wands, D. 2003, Phys. Rev. D, 67, 023503
Massardi, M., Bonaldi, A., Bonavera, L., et al. 2011, MNRAS, 415, 1597
Mennella, A., Butler, R. C., Curto, A., et al. 2011, A&A, 536, A3
Mitra, S., Rocha, G., Górski, K. M., et al. 2011, ApJS, 193, 5
Miville-Deschênes, M.-A., & Lagache, G. 2005, ApJS, 157, 302
Miville-Deschênes, M.-A., Lagache, G., Boulanger, F., & Puget, J.-L. 2007, A&A, 469, 595
Miville-Deschênes, M.-A., Ysard, N., Lavabre, A., et al. 2008, A&A, 490, 1093
Moroi, T., & Takahashi, T. 2001, Phys. Lett. B, 522, 215
Negrello, M., Clemens, M., Gonzalez-Nuevo, J., et al. 2013, MNRAS, 429, 1309
Piffaretti, R., Arnaud, M., Pratt, G. W., Pointecouteau, E., & Melin, J.-B. 2011, A&A, 534, A109
Planck Collaboration 2015, The Explanatory Supplement to the Planck 2015 results, http://wiki.cosmos.esa.int/planckpla/index.php/Main_Page (ESA)
Planck Collaboration XX. 2011, A&A, 536, A20
Planck Collaboration I. 2014, A&A, 571, A1
Planck Collaboration VI. 2014, A&A, 571, A6
Planck Collaboration IX. 2014, A&A, 571, A9
Planck Collaboration XI. 2014, A&A, 571, A11
Planck Collaboration XV. 2014, A&A, 571, A15
Planck Collaboration XX. 2014, A&A, 571, A20
Planck Collaboration XXVII. 2014, A&A, 571, A27
Planck Collaboration XXIX. 2014, A&A, 571, A29
Planck Collaboration XXX. 2014, A&A, 571, A30
Planck Collaboration I. 2016, A&A, 594, A1
Planck Collaboration II. 2016, A&A, 594, A2
Planck Collaboration III. 2016, A&A, 594, A3
Planck Collaboration IV. 2016, A&A, 594, A4
Planck Collaboration V. 2016, A&A, 594, A5
Planck Collaboration VI. 2016, A&A, 594, A6
Planck Collaboration VII. 2016, A&A, 594, A7
Planck Collaboration VIII. 2016, A&A, 594, A8
Planck Collaboration IX. 2016, A&A, 594, A9
Planck Collaboration X. 2016, A&A, 594, A10
Planck Collaboration XI. 2016, A&A, 594, A11
Planck Collaboration XII. 2016, A&A, 594, A12
Planck Collaboration XIII. 2016, A&A, 594, A13
Planck Collaboration XIV. 2016, A&A, 594, A14
Planck Collaboration XV. 2016, A&A, 594, A15
Planck Collaboration XVI. 2016, A&A, 594, A16
Planck Collaboration XVII. 2016, A&A, 594, A17
Planck Collaboration XVIII. 2016, A&A, 594, A18
Planck Collaboration XIX. 2016, A&A, 594, A19

---

[16] http://crd.lbl.gov/cmb-data

Planck Collaboration XX. 2016, A&A, 594, A20
Planck Collaboration XXI. 2016, A&A, 594, A21
Planck Collaboration XXII. 2016, A&A, 594, A22
Planck Collaboration XXIII. 2016, A&A, 594, A23
Planck Collaboration XXIV. 2016, A&A, 594, A24
Planck Collaboration XXV. 2016, A&A, 594, A25
Planck Collaboration XXVI. 2016, A&A, 594, A26
Planck Collaboration XXVII. 2016, A&A, 594, A27
Planck Collaboration XXVIII. 2016, A&A, 594, A28
Planck Collaboration Int. XV. 2014, A&A, 565, A103
Planck Collaboration Int. XVII. 2014, A&A, 566, A55
Planck Collaboration Int. XXII. 2015, A&A, 576, A107
Planck HFI Core Team 2011, A&A, 536, A4
Prezeau, G., & Reinecke, M. 2010, ApJS, 190, 267
Reinecke, M., Dolag, K., Hell, R., Bartelmann, M., & Enßlin, T. A. 2006, A&A, 445, 373
Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011, MNRAS, 418, 467
Ricci, R., Prandoni, I., Gruppioni, C., Sault, R. J., & De Zotti, G. 2004, A&A, 415, 549
Sadler, E. M., Ricci, R., Ekers, R. D., et al. 2008, MNRAS, 385, 1656
Scott, K. S., Wilson, G. W., Aretxaga, I., et al. 2012, MNRAS, 423, 575
Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709
Verde, L., Wang, L., Heavens, A. F., & Kamionkowski, M. 2000, MNRAS, 313, 141
Wang, L., & Kamionkowski, M. 2000, Phys. Rev. D, 61, 063504
Watson, R. A., Rebolo, R., Rubiño-Martín, J. A., et al. 2005, ApJ, 624, L89
Zonca, A., Franceschet, C., Battaglia, P., et al. 2009, J. Instr., 4, T12010

[1] APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/lrfu, Observatoire de Paris, Sorbonne Paris Cité, 10 rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France

[2] African Institute for Mathematical Sciences, 6−8 Melrose Road, Muizenberg, 7945 Cape Town, South Africa

[3] Agenzia Spaziale Italiana Science Data Center, via del Politecnico snc, 00133 Roma, Italy

[4] Aix Marseille Université, CNRS, LAM (Laboratoire d'Astrophysique de Marseille) UMR 7326, 13388 Marseille, France

[5] Astrophysics Group, Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge CB3 0HE, UK

[6] Astrophysics & Cosmology Research Unit, School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, 4000 Durban, South Africa

[7] Atacama Large Millimeter/submillimeter Array, ALMA Santiago Central Offices, Alonso de Cordova 3107, Vitacura, Casilla 763 0355 Santiago, Chile

[8] CITA, University of Toronto, 60 St. George St., Toronto, ON M5S 3H8, Canada

[9] CNRS, IRAP, 9 Av. colonel Roche, BP 44346, 31028 Toulouse Cedex 4, France

[10] CRANN, Trinity College, CA 91125 Dublin, Ireland

[11] California Institute of Technology, Pasadena, PA 91125 California, USA

[12] Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK

[13] Centro de Estudios de Física del Cosmos de Aragón (CEFCA), Plaza San Juan 1, planta 2, 44001 Teruel, Spain

[14] Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, 92093 California, USA

[15] Consejo Superior de Investigaciones Científicas (CSIC), 0424 Madrid, Spain

[16] DSM/Irfu/SPP, CEA-Saclay, 91191 Gif-sur-Yvette Cedex, France

[17] DTU Space, National Space Institute, Technical University of Denmark, Elektrovej 327, 2800 Kgs. Lyngby, Denmark

[18] Département de Physique Théorique, Université de Genève, 24 Quai E. Ansermet, 1211 Genève 4, Switzerland

[19] Departamento de Astrofísica, Universidad de La Laguna (ULL), 38206 La Laguna, Tenerife, Spain

[20] Departamento de Física, Universidad de Oviedo, Avda. Calvo Sotelo s/n, 33007 Oviedo, Spain

[21] Department of Astronomy and Astrophysics, University of Toronto, 50 Saint George Street, Toronto, ON M5S 3H4, Ontario, Canada

[22] Department of Astrophysics/IMAPP, Radboud University Nijmegen, PO Box 9010, 6500 GL Nijmegen, The Netherlands

[23] Department of Physics & Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, British Columbia, Canada

[24] Department of Physics and Astronomy, Dana and David Dornsife College of Letter, Arts and Sciences, University of Southern California, Los Angeles, CA 90089, USA

[25] Department of Physics and Astronomy, University College London, London WC1E 6BT, UK

[26] Department of Physics, Florida State University, Keen Physics Building, 77 Chieftan Way, Tallahassee, Florida, USA

[27] Department of Physics, Gustaf Hällströmin katu 2a, University of Helsinki, 00560 Helsinki, Finland

[28] Department of Physics, Princeton University, Princeton, CA 93106 New Jersey, USA

[29] Department of Physics, University of California, Santa Barbara, California, USA

[30] Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, Illinois, USA

[31] Dipartimento di Fisica e Astronomia G. Galilei, Università degli Studi di Padova, via Marzolo 8, 35131 Padova, Italy

[32] Dipartimento di Fisica e Scienze della Terra, Università di Ferrara, via Saragat 1, 44122 Ferrara, Italy

[33] Dipartimento di Fisica, Università La Sapienza, P.le A. Moro 2, 00133 Roma, Italy

[34] Dipartimento di Fisica, Università degli Studi di Milano, via Celoria, 16, 00133 Milano, Italy

[35] Dipartimento di Fisica, Università degli Studi di Trieste, via A. Valerio 2, 34127 Trieste, Italy

[36] Dipartimento di Matematica, Università di Roma Tor Vergata, via della Ricerca Scientifica, 1, 000138 Roma, Italy

[37] Discovery Center, Niels Bohr Institute, Blegdamsvej 17, 1165 Copenhagen, Denmark

[38] European Southern Observatory, ESO Vitacura, Alonso de Cordova 3107, Vitacura, Casilla 19001, Santiago, Chile

[39] European Space Agency, ESAC, Planck Science Office, Camino bajo del Castillo, s/n, Urbanización Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain

[40] European Space Agency, ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands

[41] Facoltà di Ingegneria, Università degli Studi e-Campus, via Isimbardi 10, 22060 Novedrate (CO), Italy

[42] Gran Sasso Science Institute, INFN, viale F. Crispi 7, 67100 L'Aquila, Italy

[43] HGSFP and University of Heidelberg, Theoretical Physics Department, Philosophenweg 16, 69120 Heidelberg, Germany

[44] Helsinki Institute of Physics, Gustaf Hällströmin katu 2, University of Helsinki, 00560 Helsinki, Finland

[45] INAF−Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, 35131 Padova, Italy

[46] INAF−Osservatorio Astronomico di Roma, via di Frascati 33, Monte Porzio Catone, Italy

[47] INAF−Osservatorio Astronomico di Trieste, via G.B. Tiepolo 11, 34127 Trieste, Italy

[48] INAF/IASF Bologna, via Gobetti 101, 40127 Bologna, Italy

[49] INAF/IASF Milano, via E. Bassini 15, 20133 Milano, Italy

[50] INFN, Sezione di Bologna, via Irnerio 46, 40126 Bologna, Italy

[51] INFN, Sezione di Roma 1, Università di Roma Sapienza, Piazzale Aldo Moro 2, 00185 Roma, Italy

[52] INFN, Sezione di Roma 2, Università di Roma Tor Vergata, via della Ricerca Scientifica 1, 00185 Roma, Italy

[53] INFN/National Institute for Nuclear Physics, via Valerio 2, 34127 Trieste, Italy

[54] IPAG: Institut de Planétologie et d'Astrophysique de Grenoble, Université Grenoble Alpes, IPAG; CNRS, IPAG, 38000 Grenoble, France

[55] IUCAA, Post Bag 4, Ganeshkhind, Pune University Campus, 411 007 Pune, India

[56] Imperial College London, Astrophysics group, Blackett Laboratory, Prince Consort Road, London, SW7 2AZ, UK

[57] Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA

[58] Institut Néel, CNRS, Université Joseph Fourier Grenoble I, 25 rue des Martyrs, 38000 Grenoble, France

[59] Institut Universitaire de France, 103 Bd Saint-Michel, 75005 Paris, France

[60] Institut d'Astrophysique Spatiale, CNRS (UMR8617) Université Paris-Sud 11, Bâtiment 121, 91898 Orsay, France

[61] Institut d'Astrophysique de Paris, CNRS (UMR7095), 98 bis Boulevard Arago, 75014 Paris, France

[62] Institut für Theoretische Teilchenphysik und Kosmologie, RWTH Aachen University, 52056 Aachen, Germany

[63] Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

[64] Institute of Theoretical Astrophysics, University of Oslo, Blindern, 0371 Oslo, Norway

[65] Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, 38205 La Laguna, Tenerife, Spain

[66] Instituto de Física de Cantabria (CSIC-Universidad de Cantabria), Avda. de los Castros s/n, Santander, Spain

[67] Istituto Nazionale di Fisica Nucleare, Sezione di Padova, via Marzolo 8, 35131 Padova, Italy

[68] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California, USA

[69] Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester, M13 9PL, UK

[70] Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

[71] Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

[72] Kazan Federal University, 18 Kremlyovskaya St., 420008 Kazan, Russia

[73] LAL, Université Paris-Sud, CNRS/IN2P3, 91898 Orsay, France

[74] LERMA, CNRS, Observatoire de Paris, 61 Avenue de l'Observatoire, 75000 Paris, France

[75] Laboratoire AIM, IRFU/Service d'Astrophysique − CEA/DSM − CNRS − Université Paris Diderot, Bât. 709, CEA-Saclay, 91191 Gif-sur-Yvette Cedex, France

[76] Laboratoire Traitement et Communication de l'Information, CNRS (UMR 5141) and Télécom ParisTech, 46 rue Barrault, 75634 Paris Cedex 13, France

[77] Laboratoire de Physique Subatomique et Cosmologie, Université Grenoble-Alpes, CNRS/IN2P3, 53 rue des Martyrs, 38026 Grenoble Cedex, France

[78] Laboratoire de Physique Théorique, Université Paris-Sud 11 & CNRS, Bâtiment 210, 91405 Orsay, France

[79] Lawrence Berkeley National Laboratory, Berkeley, 94720 California, USA

[80] Lebedev Physical Institute of the Russian Academy of Sciences, Astro Space Centre, 84/32 Profsoyuznaya st., Moscow, 117997 GSP-7, Russia

[81] Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

[82] McGill Physics, Ernest Rutherford Physics Building, McGill University, 3600 rue University, Montréal, QC, H3A 2T8, Canada

[83] National University of Ireland, Department of Experimental Physics, Maynooth, Co. Kildare, Ireland

[84] Nicolaus Copernicus Astronomical Center, Bartycka 18, 00-716 Warsaw, Poland

[85] Niels Bohr Institute, Blegdamsvej 17, 1165 Copenhagen, Denmark

[86] SISSA, Astrophysics Sector, via Bonomea 265, 34136, Trieste, Italy

[87] School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff, CF24 3AA, UK

[88] School of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD, UK

[89] Sorbonne Université-UPMC, UMR7095, Institut d'Astrophysique de Paris, 98 bis Boulevard Arago, 75014 Paris, France

[90] Space Sciences Laboratory, University of California, Berkeley, California, USA

[91] Special Astrophysical Observatory, Russian Academy of Sciences, Nizhnij Arkhyz, Zelenchukskiy region, 369167 Karachai-Cherkessian Republic, Russia

[92] Sub-Department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

[93] Theory Division, PH-TH, CERN, 1211 Geneva 23, Switzerland

[94] UPMC Univ. Paris 06, UMR7095, 98 bis Boulevard Arago, 75014 Paris, France

[95] Université de Toulouse, UPS-OMP, IRAP, 31028 Toulouse Cedex 4, France

[96] University Observatory, Ludwig Maximilian University of Munich, Scheinerstrasse 1, 81679 Munich, Germany

[97] University of Granada, Departamento de Física Teórica y del Cosmos, Facultad de Ciencias, 18010 Granada, Spain

[98] University of Granada, Instituto Carlos I de Física Teórica y Computacional, 18071 Granada, Spain

[99] Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warszawa, Poland

Table A.1. Spurious map pointing weights.

Detector	Horn 1	Horn 2	Horn 3	Horn 4	Horn 5	Horn 6	Horn 7	Horn 8	Horn 9	Horn 10	Horn 11
LFI18M	+1	0	0	0	0	0	+1	0	0	+1	+1
LFI18S	−1	0	0	0	0	0	+1	0	0	+1	+1
LFI19M	0	+1	0	0	0	0	−1	0	0	+1	+1
LFI19S	0	−1	0	0	0	0	−1	0	0	+1	+1
LFI20M	0	0	+1	0	0	0	0	+1	0	−1	+1
LFI20S	0	0	−1	0	0	0	0	+1	0	−1	+1
LFI21M	0	0	0	+1	0	0	0	−1	0	−1	+1
LFI21S	0	0	0	−1	0	0	0	−1	0	0	+1
LFI22M	0	0	0	0	+1	0	0	0	+1	0	−1
LFI22S	0	0	0	0	−1	0	0	0	+1	0	−1
LFI23M	0	0	0	0	0	+1	0	0	−1	0	−1
LFI23S	0	0	0	0	0	−1	0	0	−1	0	−1

Notes. Horns 1–6 are the physical horns that comprise two polarization-orthogonal arms. Horns 7–9 are the "virtual horns"; each consists of two physical horns, LFI18+LFI19, LFI20+LFI21, and LFI22+LFI23. Horns 10–11 expand the hierarchy to include mismatch between the virtual horns.

Appendix A: Measuring band-pass mismatch residuals through spurious-component maps

We measure bandpass mismatch by projecting the mismatch into extra degrees of freedom called "spurious" maps. The mismatch has a specific signature: it is a sky-synchronous difference between the detectors at a given frequency that does not modulate under rotation (unlike polarization). We have extended the method to measure the mismatch between individual detectors and the noise-weighted frequency average. Conventionally, the mismatch is only resolved between two polarization-orthogonal detectors that have common pointing (i.e., are co-located on the focal plane).

Mathematically we can formulate this as an extension of the pointing model. We start from the conventional polarized detector response:

$$d_t^{\mathrm{D}} = I_p + \eta^{\mathrm{D}} \left(Q_p \cos 2(\alpha_t + \psi^{\mathrm{D}}) + U_p \sin 2(\alpha_t + \psi^{\mathrm{D}}) \right) + n_t, \quad (A.1)$$

where detector D TOD sample, d_t, at time t is a linear combination of the observed intensity sky, I, at pixel p and linearly polarized (Q, U) sky that is modulated by the polarization efficiency, η^{D}, the polarization sensitive angle of the detector, ψ^{D} and the orientation of the detector α_t. In this model, the sky signal includes the azimuthally symmetric part of the beam, and all instrumental noise and systematics are contained in the noise, n. The full data vector can be written concisely as

$$\boldsymbol{d} = \mathsf{P}\boldsymbol{m} + \boldsymbol{n}, \quad (A.2)$$

where P is the pointing matrix that scans the smoothed sky map, $\boldsymbol{m}$, into the time domain according to the instrument scanning pattern.

We extend the model by assuming that the detectors are actually observing sky emission that is convolved with the individual detector bandpass, causing each detector to view a slightly different sky, I^{D}. The difference between an effective frequency average, I, and the sky seen by a particular detector, I^{D}, is the spurious (or bandpass mismatch) map, $S^{\mathrm{D}} = I - I^{\mathrm{D}}$. We can solve for the frequency-averaged (I, Q, U) sky by introducing a hierarchy of spurious maps that captures the mismatch first within pairs of detectors (horns), then between pairs of horns (quadruplets), and ultimately between larger aggregates. We call all these combinations "virtual horns". We show in Table A.1 how the spurious maps can be assigned for the *Planck* 70 GHz channel. The specific hierarchy of virtual horns is not important: it is only required that for each pair of detectors there exists exactly one spurious map (virtual or physical horn) for which they have opposite pointing weights. The total bandpass mismatch map between any detector and the frequency average can be produced as a linear combination of the spurious maps with the appropriate pointing weights.

The spurious maps extend the pointing model in Eq. (A.1) to

$$\widetilde{d_t^{\mathrm{D}}} = d_t^{\mathrm{D}} + \sum_i w_i S_i, \quad (A.3)$$

where the sum is over all spurious maps that correspond to physical and virtual horns listed in Table A.1 and the associated weights are +1, −1, or 0. These extra pointing weights are included in additional columns in the pointing matrix, P, and a full map with all the spurious components can be binned using the standard mapmaking equation:

$$\widetilde{\boldsymbol{m}} = \left(\mathsf{P}^{\mathrm{T}} \mathsf{N}^{-1} \mathsf{P} \right)^{-1} \mathsf{P}^{\mathrm{T}} \mathsf{N}^{-1} \boldsymbol{d}, \quad (A.4)$$

where N is some model of the detector noise covariance, frequently a diagonal matrix with the detector white noise variances. In practice, the additional degrees of freedom in the spurious maps can only be resolved for pixels that have sufficient coverage in crossing angles to break the degeneracy between polarization and bandpass mismatch. Even then, the resulting noise levels in the I, Q, and U maps are much higher than in the standard mapmaking case, making the $IQUS$ solutions only useful for low-resolution studies and checking other bandpass leakage correction schemes that rely on prior information such as detector bandpasses and models of the sky emission.